INTRODUCTION

TO

POVERTY ANALYSIS

World Bank Institute

August 2005

Table of Contents

Introduction

This manual presents an introductory course on poverty analysis. It covers all the basic methods related to poverty measurement and diagnosis, and applies these methods using household survey data. The topics included in this manual are: the concept of well-being, why measure poverty, how measure poverty, setting poverty lines, poverty indices and their comparisons, inequality measures, poverty profiles, the determinants of poverty, whether poverty is falling worldwide, and how poverty analysis is linked to poverty reduction policies. In order to make the material as clear as possible, the manual includes extensive illustrations from around the world.

The manual was originally developed to be used in a full-time ten-day training course of morning lectures and afternoon practice exercises, geared towards participants with at least a university-level background in science or social science. The course was designed as part of a broader training program of the World Bank Institute called the Poverty Analysis Initiative (PAI), whose objective is to improve the capacity of analysts, researchers and statisticians in less-developed countries to undertake poverty analysis, directly using data from household surveys. Among other things, this should help support the policy work related to the PRSP (Poverty Reduction Strategy Paper) process involving poverty identification, measurement, monitoring, and evaluation.

The World Bank Institute has used a draft version of the manual in training workshops in the Philippines, Cambodia and Malawi, as well as in distance courses with substantial numbers of participants in several countries in Asia (in 2002) and Africa (in 2003). The feedback from these courses has been very useful in helping us create a manual that balances rigor with accessibility and practicality. In one form or another the manual has been easily accessible on the Web for some time.

It is also possible to master the material without a formal course. This is because the manual not only sets out the tools required for undertaking basic poverty analysis with household data, using STATA software, but also includes exercises that illustrate the techniques and measures discussed in the text. Readers intending to master the material in the manual should expect to work through all the exercises in detail. In our experience, anyone who participates fully in a 10-day training course using these materials, or who works though the manual and exercises carefully, emerges well prepared to undertake relatively sophisticated poverty analyses.

Although the manual discuses all the basic elements required for poverty analysis, the coverage is necessarily somewhat selective. A more advanced course would pay more attention to the techniques used to analyze panel data, and related issues of measuring chronic vs. transient poverty and computing transition matrices. It would explain how to develop and use formal tests of changes in poverty, and would discuss econometric issues in much more depth. Such a course would consider in more detail how to target the poor and how to operationalize poverty mapping, and it would look at program and project impact evaluation, include the issues that arise in experimental design. It would also provide more detail on how shocks, such as the Asian Financial Crisis of 1997, affect poverty. It is expected that a more advanced course in poverty analysis that covers these topics, and accompanying written materials, will be developed in the near future.

This manual has been prepared under the general direction of Shahid Khandker of the World Bank Institute, and was edited by Jonathan Haughton. Much of the manual is based on a draft version written by Jonathan Haughton in 2001, supplemented with extensive materials drawn from the work of Martin Ravallion of the World Bank's Development Research Group, as well as background papers by Kevin Carey and Zeynep Orhun, and contributions from Kathleen Beegle, Celia Reyes, José Ramon ("Toots") Albert and Nidhiya Menon. The complete manual was reviewed fully by Shahid Khandker and Zeynep Orhun. The STATA exercises were prepared by Hussain Samad and Changqing Sun. A French-language version of the manual is also available, translated by Sylvie Russo.

Questions, comments and suggestions related to the manual are most welcome, and should be directed to Shahid Khandker at skhandker@worldbank.org. Our goal is to enhance local capacity in poverty analysis, and we hope that this manual represents a useful step in this direction.

Chapter 1. What is Poverty and Why Measure it?

Summary

Poverty is "pronounced deprivation in well-being." The conventional view links well-being primarily to command over commodities, so the poor are those who do not have enough income or consumption to put them above some adequate minimum threshold. This view sees poverty largely in monetary terms.

Poverty may also be tied to a specific type of consumption; thus someone might be house poor or food poor or health poor. These dimensions of poverty can often be measured directly, for instance by measuring malnutrition or literacy.

The broadest approach to well-being (and poverty) focuses on the "capability" of the individual to function in society. The poor lack key capabilities, and may have inadequate income or education, or be in poor health, or feel powerless, or lack political freedoms.

There are four reasons to measure poverty. First, to keep the poor on the agenda; if poverty were not measured, it would be easy to forget the poor. Second, one needs to be able to identify the poor if one is to be able to target interventions that aim to reduce or alleviate poverty. Third, to monitor and evaluate projects and policy interventions that are geared towards the poor. And finally, to evaluate the effectiveness of institutions whose goal is to help the poor.

To help countries think clearly and systematically about how the position of the poor may be improved, and to act in consequence, the World Bank favors the Poverty Reduction Strategy Paper (PRSP) process. Countries are expected to measure and analyze domestic poverty, and to identify and operationalize actions to reduce poverty. The PRSP process requires strong technical support. A central purpose of this manual is to impart the technical and analytical skills needed for this purpose.

Learning Objectives

After completing the chapter on *What is Poverty and Why Measure It?* you should be able to:

1. Define poverty.
2. Summarize the three main views of poverty.
3. State four justifications for measuring poverty.
4. Summarize the role of the Poverty Reduction Strategy Paper (PRSP) process.
5. Explain why technical and analytical training in poverty analysis is needed.

1.1 The concept of well-being and poverty

According to the World Bank (2000), "poverty is pronounced deprivation in well-being." This of course begs the question of what is meant by well-being.

One approach is to think of one's well-being as the command over commodities in general, so people are better off if they have a greater command over resources. In this view, the main focus is on whether households or individuals have enough resources to meet their needs. Typically poverty is then measured by comparing an individual's income or consumption with some defined threshold below which they are considered to be poor. This is the most conventional view – poverty is seen largely in monetary terms – and is the starting point for most analyses of poverty.

A second approach to well-being (and hence poverty) is to ask whether people are able to obtain a *specific* type of consumption good: do they have enough food? Or shelter? Or health care? Or education? In this view the analyst would need to go beyond the more traditional monetary measures of poverty: nutritional poverty might be measured by examining whether children are stunted or wasted; and educational poverty might be measured by asking whether someone is illiterate, or by the amount of formal schooling they have received.

Perhaps the broadest approach to well-being (and poverty) is the one articulated by Amartya Sen (1987), who argues that well-being comes from a "capability" to function in society. Thus poverty arises when people lack key capabilities, and so have inadequate income or education, or poor health, or insecurity, or low self confidence, or a sense of powerlessness, or the absence of rights such as freedom of speech. Viewed in this way, poverty is a multi-dimensional phenomenon, and less amenable to simple solutions. So, for instance, while higher average incomes will certainly help reduce poverty, these may need to be accompanied by measures to empower the poor, or insure them against risks, or to address specific weaknesses (such as inadequate availability of schools or a corrupt health service).

Poverty is related to, but distinct from, inequality and vulnerability. Inequality focuses on the distribution of attributes, such as income or consumption, across the whole population. In the context of poverty analysis, inequality requires examination if one believes that the welfare of an individual depends

on their economic position relative to others in society. Vulnerability is defined as the risk of falling into poverty in the future, even if the person is not necessarily poor now; it is often associated with the effects of "shocks" such as a drought, a drop in farm prices, or a financial crisis. Vulnerability is a key dimension of well-being since it affects individuals' behavior (in terms of investment, production patterns, and coping strategies) and the perceptions of their own situations.

The concepts, measures and analytical tools covered in this manual are mainly introduced in the context of the monetary measures of poverty, especially consumption. However, they frequently are, and should be, applied to other the dimensions of poverty.

1.2 Why measure poverty?

It takes time, energy and money to measure poverty, since it can only be done properly by gathering survey data directly from households. Why, then, do we need to go to the trouble of measuring poverty? At least four good reasons come to mind.

To keep the poor on the agenda

Perhaps the strongest justification is that provided by Ravallion (1998), who argues, "a credible measure of poverty can be a powerful instrument for focusing the attention of policy makers on the living conditions of the poor." Put another way, it is easy to ignore the poor if they are statistically invisible. The measurement of poverty is thus necessary if it is to appear on the political and economic agenda.

To target interventions, domestically and worldwide

A second reason for measuring poverty is in order to target interventions. Clearly, one cannot help the poor without at least knowing who they are. This is the purpose of a poverty profile, which sets out the major facts on poverty (and, typically, 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 is invaluable, even though it typically uses rather basic techniques such as tables and graphs. For a straightforward example, see Nicholas Prescott and Menno Pradhan, *A Poverty Profile of Cambodia* (1997).

Probably the most important operational use of the poverty profile is to support efforts to target development resources towards poorer areas. But which regions should command priority in targeting? This question can only be answered at a highly aggregate level by most survey data (like the Cambodian SESC of 1993/94 and the CSES of 1999) because of the limited number of geographic domains that were sampled. For example, in the CSES 1999, poverty is lowest in Phnom Penh, where the headcount poverty rate was 15% compared to the national poverty rate of 51%. The survey data can sometimes be combined with more detailed census data to allow for much finer geographic targeting.

A good poverty profile also makes employment targeting possible. The ability of the vast majority of households in Cambodia to escape poverty will depend on their earnings from employment. The highest poverty rate was found among people living in households headed by farmers (46% in 1993/94 in Cambodia). By contrast, households headed by someone working in the government are least likely to be poor; in these occupations the poverty rate was 20% (1993/94). This would suggest that policies that aim to reduce poverty through enhancing income-generating capabilities should be targeted towards the agricultural sector.

The relationship between poverty and education is particularly important because of the key role played by education in raising economic growth and reducing poverty. The better educated have higher incomes and thus are much less likely to be poor. Cambodians living in households with an uneducated household head are more likely to be poor, with a poverty rate of 47% in 1993/94. With higher levels of education, the likelihood of being poor falls considerably. Raising education attainment is clearly a high priority in order to improve living standards and reduce poverty.

The relationship between gender and poverty may also indicate another targeting strategy for poverty reduction. In Cambodia about 25% of the population lives in households headed by women. Perhaps surprisingly, the CSES 1999 data show that, the poverty rate was slight *lower* among female-headed households (48%) than among male-headed households (52%). In this case, targeting interventions based on the gender of the head of household would not help to distinguish the poor from the non-poor.

Targeting is also important at a world-wide level. Institutions, such as the World Bank and aid agencies, have limited resources, and would like to know how best to deploy those resources in order to combat poverty. For this, they need to know where in the world the poor are located, and this in turn requires viable information on poverty in every country. All developed countries, and about two-thirds of

developing countries, have undertaken nationally representative household surveys to collect information on consumption and/or income; in many cases, these surveys have been repeated over time.

But successful efforts to target policies and programs to help the poor also require an understanding of why they are poor. This is not simply academic curiosity: it is integral to the process of finding workable solutions and managing tradeoffs. For instance, does a tax on rice exports help the poor? We know it will favor urban residents who eat rice and will hurt rice farmers, but more information is needed before we can conclude that the policy would help the poor. Or will providing outboard motors help poor fishermen? It might simply lead to overfishing, and so be of no long-term help. And will providing sewers in slums help the poor residents, or might it worsen their lot as higher rents force them to move and provide a windfall to landowners? Questions such as these cannot be answered adequately without viable information that measures poverty.

To monitor and evaluate projects and policy interventions geared towards the poor

More generally, the third reason for measuring poverty is to be able to predict the effects of, and then evaluate, policies and programs designed to help the poor. Policies that look good on paper – new opportunities for microcredit for the poor, for instance – may, in practice, not work as well as expected. To judge the effects, one would ideally like to monitor the effects of a policy on the poor, and evaluate the outcomes in comparison with a control group. Rigorous analysis of this kind is needed both to improve the design of projects and programs, and to weed out ones that are not working.

Information on poverty is also helpful in understanding the politics of many government policies. By collecting information on households and their economic status, one can assess who uses public services and who gains from government subsidies. If programs are cut or there is retrenchment of the public sector, poverty data help inform us of the effects of these plans on the poor. Using information on poverty, one can simulate the impact of different policies. The identification of the gainers and losers goes a long way towards determining who will support, or oppose, a given policy.

To evaluate the effectiveness of institutions whose goal is to help the poor

The fourth reason for measuring poverty is to help evaluate institutions. One cannot tell if a government is doing a good job of combating poverty unless there is good information on poverty. This does not only apply to governments. "Our dream is a world free of poverty," writes the World Bank, and

its first mission statement is "to fight poverty with passion and professionalism for lasting results." The institution's success in pursuing this goal can only be judged if there are adequate measures of poverty.

1.3 Thinking Systematically: Poverty Reduction Strategy Papers (PRSP)

Measurement is necessary, but not sufficient. It is also important to think clearly and systematically about how the position of the poor may be improved, and to act in consequence.

In order to do this, the World Bank favors the Poverty Reduction Strategy Paper (PRSP) process. First introduced for Highly Indebted Poor Countries (HIPC) in 1999, this is supposed to be a country-driven policy paper setting out a long-term strategy for fighting poverty and rooted in the latest available data and analysis.

The idea is that leaders, administrators, analysts, and others from within a country should take the lead in developing a PRSP, so that the process is "owned" locally and not imposed from the outside. This begins with the measurement of poverty, followed by an analysis of its dimensions and causes. Based on this foundation, the expectation is that there will be extensive dialog about what needs to be done to reduce the number of poor. Thus, once poverty is measured and the poor are identified, the next steps in the PRSP are to choose public actions and programs that have the greatest impact on poverty, identify indicators of progress, and monitor change in a systematic manner. Poverty measurement and diagnostics are therefore central to informing policy making for poverty reduction in many countries.

The creation of a good PRSP requires strong technical support. A central purpose of this manual is to impart the technical and analytical skills needed for this purpose.

Chapter 2. 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 lowers the precision of the results, and this calls for some adjustments when analyzing the data.

The World Bank-inspired Living Standards Measurement Surveys (LSMS) feature multi-topic 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 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 OECD scale ($= 1 + 0.7 \times (N_A - 1) + N_C$) 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). But 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 module on *Measuring Poverty*, you should be able to:

6. Summarize the three steps required to measure poverty.
7. 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.
8. Describe the main problems that arise with survey data, including
    a. survey design (sampling frame/coverage, response bias),
    b. stratification, and
    c. multistage cluster sampling.
9. Explain why weighting is needed when surveys use stratified random sampling.
10. Describe and evaluate the use of equivalence scales (including the OECD scale).
11. Define consumption and income as measures of welfare, and evaluate the desirability of each in the LDC context.
12. Summarize the problems that arise in measuring income and consumption, and explain how to value durable goods, and housing services.
13. 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.
14. Argue the case that there is no ideal measure of welfare.

2.1 Steps in measuring poverty

The goal of this chapter is to set out a method for measuring poverty. There is an enormous literature on the subject, so we just set out the main practical issues, with some suggestions for further reading for those interested in pursuing the subject in more depth.

Three steps need to be taken in measuring poverty (for more discussion see Ravallion, 1998). These are:

- Defining an indicator of welfare;
- Establishing a minimum acceptable standard of that indicator to separate the poor from the non-poor (the poverty line), and;
- 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, while chapter 3 discusses the issues involved in setting a poverty line and chapter 4 deals with measuring aggregate welfare and its distribution.

2.2 Household surveys

2.2.1 Key survey issues

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 the data with care. The analyst should be aware of the following issues (see Ravallion (1999) for details):

i) *The sample frame*: The survey may represent a whole country's population, or some more narrowly defined sub-set, 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.

ii) *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.

iii) *The number of observations over time*: A single cross-section, based on one or two interviews, is the most common. *Longitudinal surveys*, in which the same households or individuals are re-surveyed over an extended period (also called panel data sets) are more difficult to do, but have been undertaken in a few countries (e.g. the Vietnam Living Standards Surveys of 1993 and 1998).

iv) *The principal living standard indicator collected*: The most common indicators used in practice are based on household consumption expenditure and household income. The most common survey used in poverty analysis is a single cross-section for a nationally representative sample, with the household as the unit of observation, and it includes data on consumption and/or income. This form of survey is cheaper per household surveyed than most alternatives, thereby allowing a larger sample than with a longitudinal or individual-based survey. A larger sample of household-level data gives greater accuracy in estimating certain population parameters, such as average consumption per capita, but can lose accuracy in estimating other variables, such as the number of under-nourished children in a population (which may require oversampling of the target group). It should not, however, be presumed that the large household consumption survey is more cost-effective for all purposes than alternatives, such as using smaller samples of individual data.

2.2.2 Common survey problems

One needs to be aware of a number of problems when interpreting household consumption or income data from a household survey.

2.2.2.1 Survey design

Even a very large sample may give biased estimates for poverty measurement if the survey is not random, or if the data extracted from it have not been corrected for possible biases, such as due to sample stratification. A random sample requires that each person in the population, or each sub-group in a stratified sample, have an equal chance of being selected.

However, the poor may not be properly represented in sample surveys; for example they may be harder to interview because they live in remote areas, or are itinerant, or live illegally in the cities and so do not appear on the rosters of the local authorities. Household surveys almost always miss one distinct sub-group of the poor: those who are homeless. Also, 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, for instance, 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 sub-groups of the poor. Or to take another example, household surveys in South Korea have typically excluded one-person households from the sample frame, which makes the results unrepresentative.

Key questions to ask about the survey are:

a) Does the sample frame (the initial listing of the population from which the sample was drawn) span the entire population?

b) Is there likely to be a response bias? This may take one of two forms – *unit non-response*, which occurs when some households do not participate in the survey, and *item non-response*, which occurs when some households do not respond fully to all the questions in the survey.

It is sometimes cost-effective deliberately to oversample some small groups (e.g. minority households in remote areas) and to undersample large and homogeneous groups. Such *stratified random sampling* – whereby different sub-groups 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 below.

2.2.2.2 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%," it would be more accurate to say something like "we are 99% confident that the true poverty rate is between 13.5% and 16.9%; our best point estimate is that it is 15.2%." Outside of academic publications, such caution is 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), in order 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 1998 (VLSS98) oversampled the sparsely-populated central highlands, and undersampled the dense and populous Red River Delta.

In cases such as this, *it is not legitimate to compute simple averages of the sample observations* (such as per capita income, for instance) in order 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 manual). The Vietnam Living Standards Survey of 1993 (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

($1,215.6), 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.

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 $943.5 (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 $1,215.6.

Table 2.1.
Illustration of why weights are needed to compute statistics based on stratified samples

	Region A	Region B	Whole country
Population (m)	2.0	8.0	10.0
True income/capita ($ p.a.)	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 =400*490	2,235,200 =1,600 * 1,397	2,431,200 =196,000 + 2,235,200
Estimated income/capita, ($ p.a.)*	490	1,397	**1,215.6** =2,431,200/2000
Case 2. Stratified sampling.			
Sample size (given initially)	1,000	1,000	2,000
Estimated total income, $	490,000 =1,000*490	1,397,000 =1,000 * 1,397	1,887,000 =490,000 + 1,397,000
Case 2a. Stratified sample, using simple average. This is incorrect, so don't do this!			
Estimated income/capita ($ p.a.)	490	1,397	**943.5** =1,887,000/2000
Case 2b. Stratified sampling, using weighted average. This is the correct approach.			
Weight (Based on population)	0.2 = 2.0/10.0	0.8 =8.0/10.0	
Estimated income/capita ($ p.a.)	490	1,397	**1,215.6** = .2*490 + .8*1,397.

Note: * 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 non-sampling error (e.g. under-reporting, poorly worded questions, etc.).

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 picked from each region (referred to as a *stratum* in the sampling context). Within each region, subregional units (towns, counties, districts, communes, etc.) 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 (e.g. 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 a couple of villages and then randomly sample 15-20 households within each chosen village. The reason for doing cluster sampling, instead of simple random sampling, is that it is cheaper. But it 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), 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 for details).

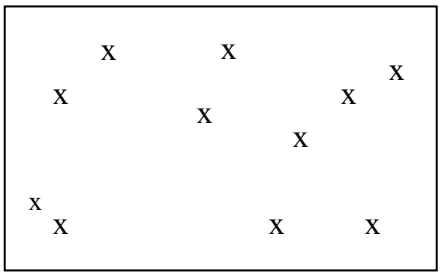

Figure 2.1.a Simple random sample

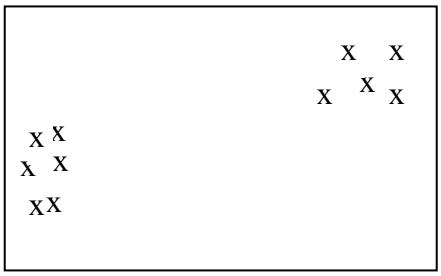

Figure 2.1.b Cluster sampling

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 many measures relate to individuals (for instance, income per capita), in which case the results need to be computed using *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.

2.2.2.3 Goods coverage and valuation

The coverage of goods and income sources in the survey should be comprehensive, including both food and non-food goods, and all income sources. Consumption should cover all monetary expenditures on goods and services consumed plus the estimated 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 income definition should include income in kind. Local market prices often provide a good guide for valuation of own-farm production or owner occupied housing.

However, whenever prices are unknown, or are an unreliable guide to reflect opportunity costs, serious valuation problems can arise. The valuation of access to public services is also difficult, and rarely done, though it is important. For transfers of in-kind goods, prevailing equivalent market prices are generally considered to be satisfactory for valuation. Non-market and durable goods present more serious problems, and there is no widely preferred method; we return to this problem in more detail below

2.2.2.4 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 (e.g. using remittances). In Less-Developed Countries, most (but not all) analysts prefer to use current consumption than current income as an indicator of living standards in poor countries, because:

i) in the short-run it reflects more accurately the resources that households control;
ii) over the long-term, it reveals information about incomes at other dates, in the past and future; and
iii) in poor countries, income is particularly difficult to measure accurately.

However, a number of factors can make current consumption a "noisy" welfare indicator. Even with ideal smoothing, consumption will still (as a rule) vary over a person's life-cycle, although this may be less of a problem in traditional societies where resource pooling within an extended family is still the norm. Another source of noise is that different households may face different constraints on their opportunities for consumption smoothing. It is generally thought that the poor are far more constrained in their ability to smooth consumption – mainly due to lack of borrowing options – than the non-poor.

2.2.2.5 Comparisons across households at similar consumption levels

Household size and demographic composition vary across households, as do the prices they face, including wage rates. As a result, it takes different resources to make ends meet for different households. In other words, at a given level of household expenditure, different households may achieve different levels of well-being: an annual income of $1,000 might suffice for a couple living in a rural area (where food and housing are cheap), but be utterly inadequate for a family of four in an urban setting.

There are a number of approaches, including equivalence scales, true cost-of-living indices, and equivalent income measures, which try to deal with this problem. The basic idea of these methods of welfare measurement is to use demand patterns to reveal consumer preferences over market goods. The consumer is assumed to maximize utility, and a utility metric is derived that is consistent with observed demand behavior, relating consumption to prices, incomes, household size, and demographic composition. The resulting measure of household utility will typically vary positively with total household expenditures, and negatively with household size and the prices faced.

The most widely-used formulation of this approach is the concept of "equivalent income", defined as 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. This gives an exact monetary measure of utility (and, indeed, it is sometimes called "money-metric utility"). Quite generally, equivalent income can be thought of as money expenditures (including the value of own production) normalized by two deflators: a suitable price index (if prices vary over the domain of the poverty comparison) and an equivalence scale (since household size and composition varies).

One of the most serious problems that arises when using equivalent income as a measure of welfare is that it does not usually include a measure of the value of access to non-market goods (e.g. public services, community characteristics), yet this varies across households. Thus two households with the same income and demographic structure may not be equally one off if one of them has access to better roads and schools, and a nicer climate.

Unfortunately there is no satisfactory solution to this particular problem, although some studies do try to include a measure of the value of at least some publicly-provided services. Information on these often comes from a separate community survey (done at the same time as the interviews, and possibly by the same interviewers), which can provide useful supplementary data on the local prices of a range of goods and local public services.

2.2.3 Key features of Living Standards Measurement (LSMS) surveys

Motivated by the need to measure poverty more accurately, the World Bank has taken a lead in the development of relatively standard, reliable household surveys, under its Living Standards

Measurement (LSMS) project. The electronic version of the books edited by Grosh and Glewwe (2000) includes sample questionnaires and detailed chapters that deal with the design and implementation of such surveys. The LSMS surveys have two key features: multi-topic questionnaires, and considerable attention to quality control. Let's consider each in more detail.

2.2.3.1 Multi-topic questionnaires

The LSMS surveys ask about a wide variety of topics, and not just 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 non-food, assets including housing, landholding and other durables, income and employment in agriculture/non-agriculture and wage/self-employment, socio-demographic variables including education, health, migration, 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.

2.2.3.2 Quality control

The LSMS surveys are distinguished by their attention to quality control. Here are some of the key features:

- Most importantly, 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 cuts down on 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, which needs to be corrected.

This concern with quality has some important implications, notably:

- 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 (e.g. province, department, county).
- The LSMS data have a fairly rapid turnaround time, with some leading to a statistical abstract (at least in draft form) within 2-6 months of the last interview.

2.3 Measuring poverty: choose an indicator of welfare

There are a number of conceptual approaches to the measurement of well-being. The most common approach is to measure economic welfare based on *household* consumption expenditure or *household* income. When divided by the number of household members, this gives a *per capita* measure of consumption expenditure or income. Of course, there are also non-monetary measures of individual welfare, which can include indicators such as infant mortality rates in the region, life expectancy, the

proportion of spending devoted to food, housing conditions, and child schooling. Well-being is a broader concept than economic welfare, which only measures a person's command over commodities.

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 derived 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:

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

where p is a vector of prices of goods and services, q is a vector of quantities of goods and services consumed, $e(.)$ is an expenditure function, x is a vector of household characteristics (e.g. number of adults, number of young children, etc.) 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. The details of this are discussed below. 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 (age, gender, job, etc).

While estimating per capita consumption might seem straightforward, there are several factors that complicate its estimation. 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 Socio-economic Survey (CSES), for example, nominal and real per capita consumption were 2,223 and 1,887 riels, respectively. However, across years the estimates 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 computed as 2,262 riels

for 1993/94, but was only 1,700 in 1999, despite economic growth during the interval; this may merely be an artifact of the different ways in which questions were asked.

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

Note: 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 & 2 in 1999 are detailed in Gibson (1999). CSES: Cambodia Socio-Economic Survey. SESC: Socio-Economic Survey of Cambodia.
Source: Gibson (1999)

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

2.3.1 Candidate 1: Income

It is tempting to measure household welfare by looking at household income. Practical problems arise immediately: what is income? and can it be measured accurately? The most generally accepted measure of income is the one formulated by Haig and Simons:

$$\text{Income} \equiv \text{consumption} + \text{change in net worth.}$$

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 (e.g. 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 Standard Survey (VLSS; undertaken in 1993 and again in 1998) 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 truly 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, etc. 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.

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**
…as % of expend.	73	70	65	60	54	**61**

Note: In 1993, exchange rate was about 10,000 dong/US$. *Source:* VLSS93

.

There are a number of reasons why income tends to be understated:

- 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 neighbors, get wind of the details.
- People may be reluctant to report income earned illegally - for instance from smuggling, or corruption, or poppy cultivation, or prostitution.

- Some parts of income are difficult to observe - e.g. the extent to which the family buffalo has risen in value.

Research based on the 1969-70 socio-economic survey in Sri Lanka estimated that wages were understated by 30%, business income by 39%, and rent, interest and dividends by 78%. 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.

2.3.2 Candidate 2: Consumption expenditure

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

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 the permanent income hypothesis, under which transitory income is saved while long-term ("permanent") income is largely consumed.

The life cycle of income and consumption is captured graphically in figure 2.2. While the available evidence does not provide strong support for this *life-cycle 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 – a month for instance – 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).

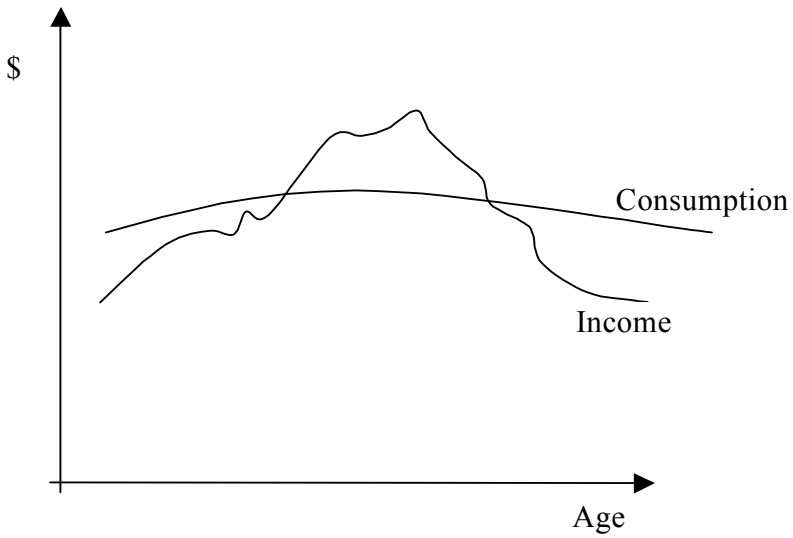

Figure 2.2 Life Cycle Hypothesis: Income and Consumption Profile over Time

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, because:

- Households tend to under-declare what they spend on luxuries (e.g. alcohol, cakes) or illicit items (drugs, prostitution). For instance, the amount that households said they spent on alcohol, according to the 1972-73 household budget survey in the US, was just half the amount that companies said they sold!
- Questions matter. According to VLSS93, Vietnamese households devoted 1.7% of their expenditure to tobacco; the VLSS98 figures showed that this had risen to 3%. An increase of this magnitude is simply not plausible, 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.

2.3.2.1 Measuring durable goods

In measuring poverty it might be argued that only food, the ultimate basic need (which anyway constitutes three quarters 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 (clothing, shelter, etc.). 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 (housing, pots and pans, etc.).

The problem here is that durable goods, such as bicycles and TVs, are bought at a point in time, and then consumed (i.e. eaten up and destroyed) 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%) during the year. Thus the true cost of the watch during the year was $8.50.

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. This suffices to compute the current consumption of the durable item, as the illustration in the following box 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 (ceteris paribus).

Box: Calculating the value of durable goods consumption - an illustration.

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

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

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

Note that 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 TV is worth 1m dong now, is expected to depreciate by 10% per annum, and the real interest rate is 3%, then the imputed consumption of the durable good is measured as 1m × (10% + 3%) = 0.13m 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.

2.3.2.2 Measure the value of housing services

If you own your house (or apartment), it provides housing services, which 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.

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 sq. ft. of floor space), the year in which it was built, the type of roof, whether there is running water, etc. This gives

Rent = f(area, running water, year built, type of roof, location, number of bathrooms, …)

This equation is then 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.[1] 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.[2]

In the case of Vietnam there is a problem with this approach: almost nobody rents housing! And 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. 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 rental value of housing was assumed to be 3 percent of the capital value of the housing. This is a somewhat arbitrary procedure, but the 3 percent is almost certainly too low.

2.3.2.3 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 (if ever) 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.

2.3.2.4 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

[1] 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.

[2] 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.

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 single individuals.

Example. For example, suppose we have a household with 2 members and monthly expenditure of $150 total. We would then assign each individual $75 as their monthly per capita expenditure. If we have another household with 3 members, it would appear that each member is worse off, with only $50 per capita per month. However, suppose we know that the 2-person household contains two adult males aged 35 whereas the second household contains 1 adult female and 2 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 measures the number of adult males (typically) to which that household is deemed to be equivalent. So 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 to even 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.

There are two possible solutions to this problem: either pick a scale that seems reasonable on the grounds that even a bad equivalence scale is better than none at all, or try to estimate a scale typically based on observed consumption behavior from household surveys. Often the equivalence scales are based on the different calorie needs of individuals of different ages.

OECD scale

Commonly used is the "OECD scale," which may be written as

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

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, etc.) of children. Osberg and Xu (1999) use the OECD scale in their study of poverty in Canada. 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 living standards measurement surveys that were undertaken in Ghana, Peru and the Côte d'Ivoire:

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.

It is also possible to estimate (econometrically) an equivalence scale, essentially by looking at how aggregate household consumption of various goods during some survey period tends to vary with household size and composition, although Deaton and Zaidi (1998) argue, "there are so far no satisfactory methods for estimating economies of scale."

A common method is to construct a demand model in which the budget share devoted to food consumption of each household is regressed on the total consumption per person. Deaton (1997) gives an example using Engel's method with household expenditure survey data from India and Pakistan. Specifically, household food share is regressed on per capita expenditure, household size, and household composition variables such as the ratio of adults and ratios of children at different ages. The equivalence scales – here the ratio of costs of a couple with a child to a couple without children – can then be calculated with the estimated coefficients. They are displayed in table 2.4:

Table 2.4: Equivalence scales using Engel's method

Age	Maharashtra, India	Pakistan
0-4	1.24	1.28
5-9	1.28	1.36
10-14	1.30	1.38
15-54	1.34	1.42

Note: Reproduced from Deaton (1997) table 4.6. Numbers show cost of 2 adults plus 1 person of the age shown, relative to a childless couple.

The numbers show the estimated costs of a family of two adults plus one additional person of various ages calculated relative to the costs of a childless couple. So, for example, a child between 0 and 4 years is equivalent to 0.24 of a couple, or 0.48 of an adult. As the age of the additional member rises, the extra costs associated with the child rise. We can compare these estimates with the last row, which shows the equivalence scale when an additional adult is added to the household. An additional adult costs 34% more for the couple, or incurs 68% of the cost of one member of the couple. So, by these calculations, these households experience economies of scale to additional adults, plus younger members are not equivalent to adults in terms of costs.

Unfortunately, there are a number of problems with this method (see Ravallion, 1994 and Deaton 1997 for details). Consider the following example from Ravallion (1994), where there are two hypothetical households as described in table 2.5.

Table 2.5: Consumption within two hypothetical households						
	Male adult	Female adult	First child	Second child	Per person	Per equivalent male adult
Household A	40	20	10	10	20	29.6
Household B	25	-	-	-	25	25
Source: Adapted from Ravallion (1994). Uses the "OECD scale": $AE = 1 + 0.7(N_{adults} - 1) + 0.5N_{children}$.						

In this example, four persons live in household A but just one in household B. The government can make a transfer to the household that is deemed to be the poorest, but it cannot observe the distribution of consumption within the households. All the government knows is the aggregate expenditure and the household composition. In this case, ***which of the two households should have priority for assistance?***

Household A has lower consumption per capita and so looks worse off. But using equivalence scales as calculated here, household B would have priority in receiving assistance. This example demonstrates two points. First, while observable consumption behavior is important information, assumptions about unobservables (e.g. how the aggregate is split within the household) will be required. Second, assumptions in computing consumption for individuals using household data can have considerable bearing on policy choices.

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 2.6.

Table 2.6. ***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 under-reported. • May be affected by short-term fluctuations (e.g. the seasonal pattern of agriculture). • Some parts of income are hard to observe (e.g. 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 standard of living. • Smoothes 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 (e.g. via borrowing, social networks). • Consumption choices made by households may be misleading (e.g. if a rich household chooses to live simply, that does 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.
Source: Based on Albert, 2004.	

2.3.3 Candidate 3. Other measures of household welfare

Even if they were measured perfectly, neither income nor expenditure would be a perfect 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.

There are other possible measures of well-being. Among the more compelling are:

- *Calories consumed per person per day.* If one accepts the 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. 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 amount of calories per person, as this will depend on the age, gender, and working activities of the individual.

Figure 2.3. Engel curve: food spending rises less quickly than income

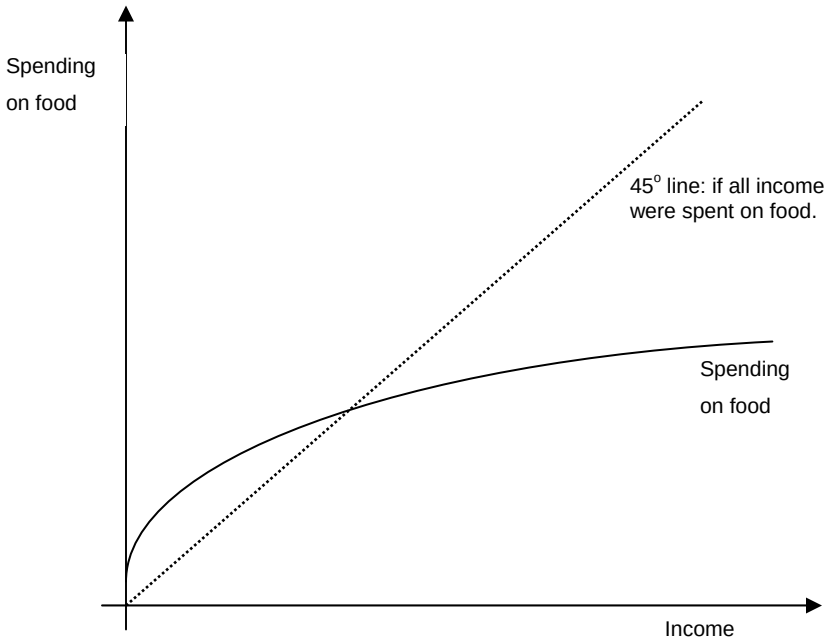

- *Food consumption as a fraction of total expenditure.* Over a century ago Ernst Engel observed, in Germany, that as household income per capita rises, spending on food rises too, ***but less quickly***. This relationship is shown in figure 2.3. As a result, the proportion of expenditure devoted to 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% 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).

- *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. 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). However, there is one further point about these measures: by some accounts, the use of child anthropometric measures to indicate nutritional need is questionable when broader concepts of well-being are invoked. For example, it has been found that seemingly satisfactory physical growth rates in children are sometimes maintained at low food-energy intake levels by not playing. That is clearly a serious food-related deprivation for any child.

- *Anthropological method.* Close observation at the household level over an extended period can provide useful supplementary information on living standards in small samples. However, this is unlikely to be a feasible method for national poverty measurement and comparisons. Lanjouw and Stern (1991) used subjective assessments of poverty in a north Indian village, based on classifying households into seven groups (very poor, poor, modest, secure, prosperous, rich and very rich) on the basis of observations and discussion with villages over that year.

  An issue of concern about this method is clearly its objectivity. The investigator may be working on the basis of an overly stylized characterization of poverty. For example, the poor in village India are widely assumed to be landless and underemployed. From the poverty profiles given by Lanjouw and Stern (1991) we find that being a landless agricultural laborer in their surveyed village is virtually a sufficient condition for being deemed poor. By their anthropological method, 99% of such households are deemed poor, though this is only so for 54% when their measurement of permanent income is used. It is clear that the perception of poverty is much more strongly linked to landlessness than income data suggest. But it is far from clear which data are telling us the most about the reality of poverty.

When one is looking at a community (e.g. 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–and nor does consumption per capita 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 2.7, which refer to eleven different

countries. How countries are ranked in terms of living standards clearly depends on which measure or indicator is considered.

Table 2.7: 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 (1992-1998)	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	na	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)					

In sum, there is no ideal measure of well-being. The implication is simple: 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.

Selected further reading:

Angus Deaton and Salman Zaidi. 1999. *Guidelines for Constructing Consumption Aggregates For Welfare Analysis*. Available at http://www.wws.princeton.edu/%7Erpds/downloads/deaton_zaidi_consumption.pdf [Accessed May 13, 2004]. Subsequently issued in 2002 as Living Standards Measurement Study Working Paper: 135. v. 104, pp. xi, Washington, D.C.: The World Bank.

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.

Margaret Grosh and Paul Glewwe. 1998. "The World Bank's Living Standards Measurement Study Household Surveys," *Journal of Economic Perspectives*, 12(1): 187-196.

Margaret Grosh and Paul Glewwe (eds.). 2000. Designing Household Survey Questionnaires for Developing Countries: Lessons from Fifteen Years of Living Standard Measurement Study. 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.

Margaret Grosh and Paul Glewwe. 1995. *A Guide to Living Standards Measurement Study Surveys and Their Data Sets*, Living Standards Measurement Study Working Paper No. 120, World Bank.

General Statistical Office (Vietnam). *Vietnam Living Standards Survey 1997-1998*, Statistical Publishing House, Hanoi 2000.

World Bank. *Vietnam: Attacking Poverty*, Hanoi, 1999.

Chapter 3. Poverty Lines

Summary

The poor are those whose expenditure (or income) falls below a poverty line. This chapter explains how poverty lines are constructed and discusses the strengths and weaknesses of defining poverty lines based on three methods: the cost of basic needs, food energy intake, and subjective evaluations. The construction of a poverty line is the most difficult step in the practical measurement of poverty.

Most commonly used is the ***cost of basic needs approach***. It first estimates the cost of acquiring enough food for adequate nutrition – usually 2,100 Calories per person per day – and then adds on the cost of other essentials such as clothing and shelter. When price information is lacking, one may use the ***food energy intake method***, which graphs expenditure (or income) per capita against food consumption (in Calories per person per day); from this one may determine the expenditure (or income) level at which a household acquires enough food. ***Subjective poverty lines*** are based on asking people what minimum income level is needed in order to just make ends meet.

An absolute poverty line remains fixed over time – adjusted only for inflation – as in the United States. It allows one to track the evolution of poverty over time, and is also useful when evaluating the effects of policies and programs on the incidence of poverty. However, in most countries poverty lines are revised from time to time; these allow one to measure relative but not absolute poverty, but typically reflect the evolution of social consensus about what constitutes poverty.

The choice of poverty line depends on the use to which it will be put: thus for international comparisons the $1/day standard is helpful, while for targeting the poor a relative poverty line suffices. The appropriate choice of poverty line is a matter of judgment, and will therefore vary from country to country.

Learning Objectives

After completing the module on poverty lines, you should be able to:

15. Explain what a poverty line is, why it is needed, and how countries adjust their poverty lines over time.
16. Distinguish between absolute and relative poverty lines, and identify the conditions under which one might be preferred to the other.
17. Identify the steps required to construct a poverty line using the Cost of Basic Needs method, and justify the choices made at each step.
18. Show how to construct a poverty line using the Food Energy Intake method, and explain the serious weaknesses of this method.
19. Explain how subjective poverty lines are constructed and critically appraise their usefulness.
20. Construct a poverty line using real survey data, using
    a. The Cost of Basic Needs method.
    b. The Food Energy Intake method.

Note: This chapter includes embedded questions, which you are encouraged to tackle as you read the text.

3.1 How to define a poverty line

Let us assume that we have chosen a measure of household well-being - say consumption expenditure. The next step is to choose a poverty line. Households whose consumption expenditure falls below this line are considered to be poor.

Following common practice, the poor are defined as those who lack command over basic consumption needs, including food and non-food components. Thus, the poverty line, thus, is obtained by specifying a consumption bundle considered adequate for basic consumption needs and then by estimating the cost of these basic needs. In other words, the poverty line may be thought of as the minimum expenditure 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 determine whether that amount places the household "in poverty", or defines the household as "poor". The threshold that we use for this is the poverty line. The poverty line defines the level of consumption (or income) needed for a household to escape poverty.

It is sometimes argued that the notion of a poverty line implies a distinct "turning point" in the welfare function. That is, by rising from just below to just above the poverty line, households (and individuals therein) move from considerable misery to an adequate minimum amount of well-being. However, given that well-being follows a continuum, and given how arbitrary the choice of poverty line is, the notion of such a "turning point" is not very compelling.

A corollary is that it may make sense to define more than one poverty line. For example, one common approach is to define one poverty line that marks households that are "poor," and another lower level that indicates those that are "extremely poor." Another approach is to construct a "food poverty line," which is based on some notion of minimum amount of money a household needs to purchase some basic-needs food bundle and nothing more. If the cost of basic non-food needs is estimated, then the food poverty line added to the non-food needs will equal the overall poverty line.

1.	A poverty line is:	
	A	The minimum expenditure required to fulfill basic needs.
	B	The threshold consumption needed for a household to escape poverty.
	C	Somewhat arbitrary because the line between poor and non poor can be hard to define.
	D	All of the above

More formally, the poverty line for a household, z_i, may be defined as the minimum spending/consumption (or income, or other measure) needed to achieve at least the minimum utility level u_z, given the level of prices (p) and the demographic characteristics of the household (x), so:

(3.1) $$z_i = e(p, x, u_Z)$$

In practice we cannot measure u_z, or even $e(.)$, and so a more pragmatic approach is needed.

There are two approaches. One is to compute a poverty line for each household, adjusting it from household to household to take into account differences in the prices they face and their demographic composition. For example, a small household in a rural area may face low housing costs and relatively modest food prices. Thus, their z_i may be low compared to a large household living in a city where housing is more expensive and food prices are perhaps higher. This gives different poverty lines for each household.

A second, and more widely used, approach is to construct one *per capita* poverty line for all individuals, but to adjust *per capita* y_i for differences in prices and household composition. The adjusted *per capita* y_i is then compared with the one poverty line to determine if the individual is living below the poverty line. With this approach, it is easier to talk of "the poverty line" and present it as a single number.

The approach taken for Cambodia in 1999 is somewhere between these two extremes. Separate poverty lines were constructed for each of three major "regions", based on the prices prevailing in those areas; whether a household in any given region is poor is then determined by comparing its expenditure per capita with the appropriate regional poverty line. These poverty lines are shown in table 3.1, based on Gibson's (1999) poverty profile of Cambodia using the CSES 1999 data, and Prescott and Pradhan's (1997) profile using the SESC 1993/94 data. We discuss the construction of these poverty lines in more detail below.

Table 3.1: Summary of Cambodia poverty lines				
	1993/94 SESC		1999 CSES	
	Food Poverty Line	Poverty Line	Food Poverty Line	Poverty Line
	(riels per person per day)			
Phnom Penh	1185	1578	1737	2470
Other Urban	995	1264	1583	2093
Rural	881	1117	1379	1777
Source: Prescott and Pradhan (1997); Gibson (1999). Average exchange rate was 2,617 riels/USD in 1993-94 and 3,808 riels/USD in 1999.				

As shown in table 3.1, poverty lines Phnom Penh, the capital of Cambodia, are higher than other areas. This is consistent with experience in other countries. For example, in Vietnam, Duong and Trinh (1999) note that the World Bank concluded that households would need to spend at least 1,071,000 dong (about US$81) per person in 1998 to be out of poverty. However, for urban areas, the amount was estimated to be 1,342,000 dong ($101); in rural areas it was just 1,054,000 dong ($79). This reflects that fact that costs are higher in cities.

Over time, we expect nominal poverty lines to change for a population. This is due to two factors. First, poverty lines reflect the costs of purchasing food and non-food items. As prices rise – inflation is typical – nominal poverty lines increase. This is what underlies the rising nominal poverty lines in Cambodia, shown in Table 3.1. It is also reflected in the poverty line for Thailand, shown in Table 3.2.

Table 3.2: Average poverty line of Thailand	
Year	Poverty Line, baht/month
1988	473
1990	522
1992	600
1994	636
1996	737
1998	878
1999	886
Note: All values are in Baht per person per month. *Source*: Kakwani, based on Thailand Socio-Economic Survey conducted by the National Statistical Office	

Second, the poverty line could change if the real poverty threshold were revised over time. This raises the question of whether we should look at relative, or absolute, poverty lines. We now consider each in turn.

2.	In measuring poverty in Cambodia, researchers used	
	A	One poverty line for the country, and adjusted household spending for price differences.
	B	Separate poverty lines for each individual.
	C	Separate poverty lines for each household.
	D	Separate poverty lines for each major region.

3.1.1 Relative poverty

Sometimes we are interested in focusing on the poorest segment (e.g. a fifth, or two-fifths) of the population; these are the relatively poor. When defined in this way, it is a truism that "the poor are always with us." It is often helpful to have a measure such as this in order to target programs that are geared to helping the poor.

In practice, rich countries have higher poverty lines than do poor countries, as shown clearly in figure 3.1 (from Ravallion 1998, p.26, based on an earlier paper by Ravallion, Datt and van de Walle). This explains why, for instance, the official poverty rate in the early 1990s was close to 15% in the United States and also close to 15% in (much poorer) Indonesia. Many of those counted as poor in the U.S. would be considered to be comfortably well off by Indonesian standards.

As countries become better off, they have a tendency to revise the poverty line upwards – with the notable exception of the United States, where the line has (in principle) remained unchanged for four decades. For instance, the European Union typically defines the poor as those whose per capita incomes fall below 50% of the median. As the median income rises, so does the poverty line.

Figure 5: Poverty lines across countries

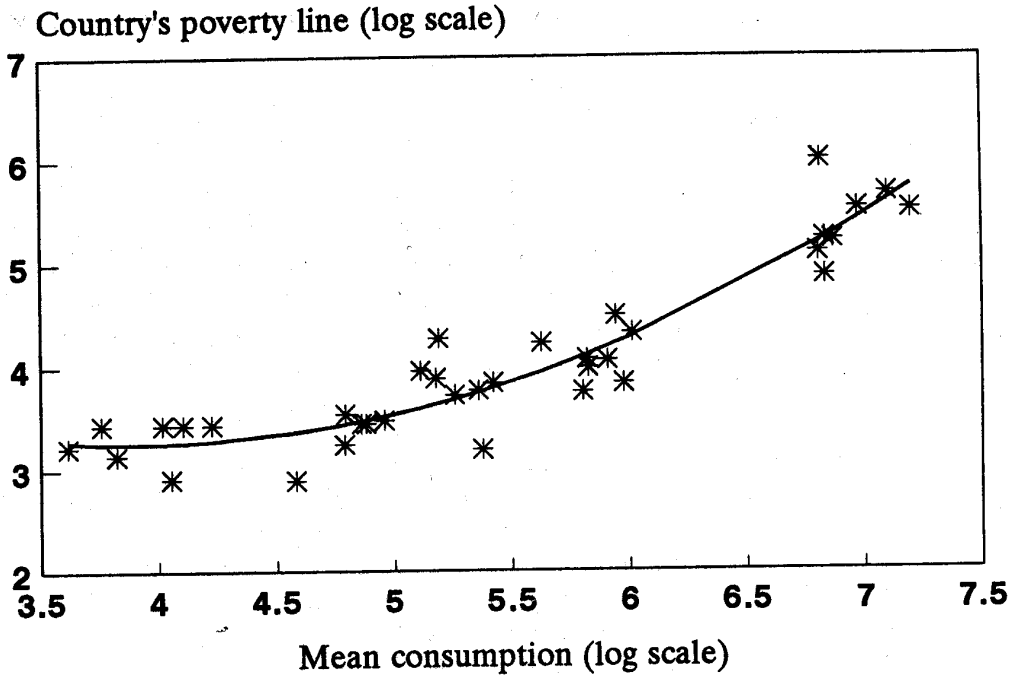

Note: Each point is one country. Both poverty line and mean consumption are measured at purchasing power parity.
Source: Ravallion, Datt and van de Walle (1991)

Figure 3.1. Poverty lines across countries.

Based on a sample of 36 countries, Ravallion, Datt and van de Walle (1991) estimated the following relationship:

(3.2) $$\ln z_i = 6.704 - 1.773 \ln(C/cap) + 0.228[\ln(C/cap)]^2 + v_i$$
t=5.1 t=-3.6 t=5.1

where R^2=0.89; all three coefficients are statistically significant at the 1% level or better. They found that at the mean value of per capita consumption (which they measured in purchasing power parity terms), the elasticity of the official poverty line (z_i) with respect to consumption per capita (C/cap) was 0.71. This means that if per capita consumption were to rise 10%, then the official poverty line would rise 7.1% on average. But the non-linear relationship implies that the elasticity of the poverty line with respect to consumption per capita was close to 0 in low-income countries, and was almost 1 in high-income countries.

To the extent that one's goal is to identify and target today's poor, then a relative poverty line is appropriate, and needs to be tailored to the overall level of development of the country. 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 (Haughton 2000), but would be of little relevance in the United States where almost nobody would be poor by this standard.

3.	According to Ravallion et al., as countries become richer, they adjust their real poverty lines upwards	
	A	A little, if they are poor, and a lot if they are rich.
	B	A little, if they are rich, and a lot if they are poor.
	C	To maintain poverty at 27%.
	D	To adjust for inflation.

3.1.2 Absolute poverty

An absolute poverty line is "fixed in terms of the standard of living it commands over the domain of poverty comparisons." In plain English, the poverty line is set so that it represents the same purchasing power year after year. For example, the United States poverty line does not change over time (except to adjust for inflation), so that the poverty rate today may be compared with the poverty rate of a decade ago, knowing that the definition of what constitutes poverty has not changed.

An absolute poverty line is essential if one is trying to judge the effect of anti-poverty policies over time, or to estimate the impact of a project (e.g. microcredit) on poverty. Legitimate comparisons of poverty rates between one country and another can only be made if the same absolute poverty line is used in both countries. Thus, the World Bank needs absolute poverty lines in order to be able to compare poverty rates across countries, which in turn is useful in determining where to channel resources, and also in assessing progress in the war on poverty. It commonly uses two measures: a) an estimated 1.1 billion people worldwide lived on less than one dollar a day in 2001 (see box for details), and b) 2.7 billion people worldwide lived on less than two dollars a day in the same year. These are absolute poverty lines. There is a vigorous controversy about whether world poverty is indeed falling – this issue is addressed more completely in chapter 10. In this context, the focus is also on absolute poverty.

Box: The "dollar a day" standard

Cross-country comparisons of poverty rates are notoriously difficult (see chapter 10), but the World Bank has tried to get around this problem by computing the proportion of the population in different countries living on less than "one dollar" per capita per day; the original line referred to $1/capita in 1985 US dollars, but was revised by Chen and Ravallion (2000) to $1.08 in 1993 US dollars (worth $1.31 in 2004 US prices). The numbers shown below suggest that the poverty rate in Vietnam (computed by Haughton 2000) compares favorably with that of India, but lags behind (more affluent) China and Indonesia. One possible lesson that may be drawn from these numbers is that the easy gains in poverty reduction in Vietnam are probably over, and the country will have difficulty reducing its poverty rate substantially in the decade ahead, even if economic growth continues at its current relatively robust rate of between 6% and 8% annually.

	% of population living on less than $1/day			% of population living on less than $1/day	
Vietnam	45	1993	Indonesia	8	1996
	27	1998	Nigeria	31	1992-93
China	22	1995	Philippines	27	1994
India	47	1994			

Sources: World Bank. 1999c. *Entering the 21st Century: World Development Report 1999/2000*. Washington DC. Haughton 2000.

4.	An absolute poverty line is needed for all of the following *except:*	
	A	To make international comparisons of poverty rates.
	B	To evaluate the effects of projects, such as irrigation investments, on poverty.
	C	To target anti-poverty measures to the poorest quintile of the population.
	D	To measure the success of government policies in combating poverty.

3.2 Issues in choosing an absolute poverty line

3.2.1 Decide the standard of living

An important conceptual problem arises when working with absolute poverty lines, which is the issue of what is meant by "the standard of living" (Ravallion, 1998, on which much of this discussion is based).

In practice, almost all absolute poverty lines are set in terms of the cost of buying a basket of goods (the "commodity-based poverty line," which we denote by z). If we assume that

(3.3) $$u = f(y),$$

which says that utility or "standard of living" (u) depends on income or expenditure (y), then

(3.4) $$y = f^{-1}(u).$$

This says that for any given level of utility, there is some income (or expenditure) level that is needed to achieve it. If u_z is the utility that just suffices to avoid being poor, then

(3.5) $$z = f^{-1}(u_z).$$

In other words, given a poverty line that is absolute in the space of welfare (i.e. gives u_z) there is a corresponding absolute commodity-based poverty line.

But suppose we make a different, but equally plausible assumption, which is that utilities are interdependent. My well-being may depend not just on what I consume, but also on how my consumption stacks up against that of the rest of society. Thus, 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 U.S., it may feel very poor. We may capture this idea by assuming

(3.6) $$u = g(y, \frac{y}{\bar{y}}),$$

where $\bar{y}$ is the mean income in the society. In this case

(3.7) $$u_z = g(z, \frac{z}{\bar{y}})$$

and so, making the standard assumption of invertibility,

(3.8) $$z = g^{-1}(\bar{y}, u_z).$$

This means that for a poverty line to be absolute in the space of welfare (i.e. to yield u_z), the commodity-based poverty line (i.e. z) may have to rise as $\bar{y}$ rises. The commodity-based poverty line would then look more like a relative poverty line! However, in what follows, we simplify the analysis by assuming that utilities are not interdependent, and so the commodity based poverty line is given in absolute terms.

5.	Is the following statement true or false? If my wellbeing depends on where I stand relative to others, then the dollar absolute poverty line needs to change as a country becomes richer.	
	True	False

3.2.2 Decide u_z *and g(.)*

Even if we assume that the commodity-based poverty line remains constant, we are still left with two problems.

a) *The Referencing problem*. What is the appropriate value of u_z – i.e. the utility of the poverty line? The choice is of course arbitrary, but "a degree of consensus about the choice of the reference utility level in a specific society may well be crucial to mobilizing resources for fighting poverty" (Ravallion, 1998, p.6).

b) *The Identification problem*. Given u_z, what is the correct value of z – i.e. of the commodity value of the poverty line. This problem arises both because the size and demographic composition of households vary – an issue we raised in the discussion of equivalence scales in chapter 2 – but also because "the view that we can measure welfare by looking solely at demand behavior is untenable" (Ravallion, 1998, p.7).

The implication is that external information and judgments will be required to answer the referencing and identification problems, and hence to determine the absolute poverty line in practice. But how is this to be done in practice?

Table 3.3 presents absolute and relative poverty headcount rates for different regions in the world. How regions compare with each other depends on which poverty measure is used. For example, by the absolute measure of less than US $1 a day, Sub-Saharan Africa has the highest portion of the population

living in poverty. On the other hand, countries in Latin America and the Caribbean have the highest portion of their population living below one-third the average national consumption; in effect, these are the most unequal societies, an issue that is addressed directly in chapter 6.

In passing we might note that an absolute poverty line is best thought of as one that is fixed in terms of living standards, and *fixed over the entire domain of the poverty comparison* (Ravallion); the domain could be a region or country, or the whole world. Thus absolute poverty comparisons will deem two persons at the same standard of living to both be either "poor" or "not poor" irrespective of the time or place being considered, or with or without some policy change, *within the relevant domain*. However, depending on the purpose of the comparison, the relevant domain may vary. For example, a global comparison of absolute consumption poverty may entail using a poverty line (e.g. $1 consumption per capita per day) that is low by the standards of rich countries. If, however, one is trying to form a poverty profile for one country only, the choice of an absolute poverty line should be appropriate to that country (e.g. a poverty line of $1 per day might be appropriate in Vietnam, and $20 per day might be suitable in the United States). Judgments of what constitutes a reasonable absolute poverty line must first specify the domain of comparisons, and recognize that the answer may change if the domain changes.

Table 3.3: Absolute and Relative Poverty rates

	Share of population living on less than $1 per day (in 1998)	Share of the population living on less than one-third of average national consumption for 1993 (in 1998)
East Asia and Pacific	15.3	19.6
East Asia and Pacific excluding China	11.3	24.6
Europe and Central Asia	5.1	25.6
Latin America and the Caribbean	15.6	51.4
Middle East and North Africa	1.9	10.8
South Asia	40.0	40.2
Sub-Saharan Africa	46.3	50.5
Total	**24.0**	**32.1**
Total excluding China	26.2	37.0
Source: World Bank (2000)		

6.	The poverty line will vary depending on the domain of comparison because of	
	A	The referencing problem.
	B	The identification problem.
	C	The purpose of the comparison.
	D	The $1/day standard is too low.

3.3 Solution A: objective poverty lines.

How then are we to determine poverty lines? One possibility is to pick an "objective" poverty line. The key idea here is that the poverty line should be set at a level that enables individuals to achieve certain capabilities, including a healthy and active life and full participation in society. In practice this almost certainly would imply that the commodity-based poverty line would rise as a country becomes more affluent, because the minimum resources needed to participate fully in society probably rise over time. In Sen's rather dense prose (Sen 1983, p.168), "an absolute approach in the space of capabilities translates into a relative approach in the space of commodities."

A common, and fairly satisfactory, way of approaching capabilities is to begin with nutritional requirements. The commonest way of making this operational is the Cost-Of-Basic Needs (CBN) approach, while the Food Energy Intake (FEI) method has been suggested as an alternative when the data available are more limited.

3.3.1 The Cost-of-Basic-Needs method:

The most satisfactory approach to building up a poverty line, while remaining in the spirit of trying to ensure that the line covers basic needs, proceeds as follows:

- Stipulate a consumption bundle that is deemed to be adequate, with both food and non-food components; and
- Estimate the cost of the bundle for each subgroup (urban/rural, each region, etc.).

This is essentially the approach taken by Seebohm Rowntree in his seminal study of poverty in York, undertaken in 1936. Note that although we begin with a consumption bundle – so much food, so much housing space, so much electricity, etc. – the poverty line is measured in money. We are therefore not insisting that each basic need be met by each person, only that it *could* be met. Operationally, the steps to follow are these:

- Pick a nutritional requirement for good health, such as 2,100 Calories per person per day. This standard is widely used, and has been proposed by the Food and Agricultural Organization of the United Nations.
- Estimate the cost of meeting this food energy requirement, using a diet that reflects the habits of households near the poverty line (e.g. those in the lowest, or second-lowest, quintile of the

income distribution; or those consuming between 2,000 and 2,200 Calories). This may not be easy if diets vary widely across the country. Call this food component z^F.

- Add a non-food component (z^{NF}). There is a lot of disagreement about how to do this; we offer some more thoughts on this issue below.
- Then the basic needs poverty line is given by

(3.12) $$z^{BN} = z^F + z^{NF}$$

7.	Is the following statement true, false or uncertain? The Cost of Basic Needs approach requires that households meet their basic needs of food and essential non-food spending.		
	True	False	Uncertain

Box. The US poverty line

In 1963 and 1964, Mollie Orshansky of the U.S. Social Security Administration computed the cost of an 'adequate' amount of food intake, to get z^F. She then multiplied this number by 3 to get z^{BN}. Why? Because at the time, the average food share for all consumers in the United States was 1/3. This line is still used, updated regularly for price changes.
Source: Dalaker and Naifeh (1998).

To illustrate how this might work, suppose, following common practice, that we use a food energy threshold of 2,100 Calories per day.[3] Suppose that there are only three foodstuffs: rice, corn and eggs. For this hypothetical example, imagine that table 3.4 shows the expenditure on each item, and the amount consumed by a household in the second (from bottom) quintile; since such a household consumes, we suppose, just 2,000 Calories per day, the figures here have to be grossed up to give the cost of purchasing 2,100 Calories. In this example the cost comes to 105 pesos per day.

Table 3.4: Illustration of Construction of Cost of Food Component of Poverty Line

	Expenditure per day (pesos)	Calories	Calories, Adjusted to give 2,100 Calories	Expenditure, adjusted to cover 2,100 Calories
Rice	60	1,400	1,470	63
Corn	20	400	420	21
Eggs	20	200	210	21
Total	**100**	**2,000**	**2,100**	**105**

[3] We use the convention that 1 Calorie is equivalent to 1,000 calories.

The choice of which diet to use when estimating the cost of obtaining 2,100 Calories is not a trivial one, a point emphasized in the context of Indonesia by Pradhan et al. (2000)[4]. To illustrate, consider the information in Table 3.5, drawn from the Vietnam Living Standards Survey of 1992-93. Households in the poorest quintile paid 0.68 dong per Calorie; those in the richest expenditure quintile paid almost twice as much (1.38 dong/Calorie). Depending on which cost/Calorie one uses, the poverty line could vary widely.

Table 3.5: Food consumption by expenditure quintile, Vietnam, 1992-93.

Quintile	Expenditure per capita, '000 dong/year	% of expenditure devoted to food	Calories per capita per day	Dong per Calorie
Lowest	562	70	1,591	0.68
Low-mid	821	65	1,855	0.79
Middle	1,075	60	2,020	0.87
Mid-upper	1,467	54	2,160	1.00
Upper	2,939	47	2,751	1.38

Source: Vietnam Living Standards Survey 1992-93

8.	In constructing a Cost of Basic Needs poverty line in Vietnam, the poverty line will be
A	Lower if the food price of the lowest quintile is used.
B	Higher if one uses the Calorie per capita level of the lowest quintile.
C	Lower if one uses the percentage of spending on non-food from the top expenditure quintile.
D	Higher if one uses a threshold of 2,020 Calories per capita per day.

An Application. In practice, researchers in this case used the price of food for households in the middle quintile, on the grounds that those households were close to the poverty line because they were consuming almost 2,100 Calories per year. The food expenditure of the middle quintile, grossed up to pay for 2,100 Calories, came to 750,228 dong per capita in 1993; the non-food expenditure of this same group of households was taken to be adequate for those at the poverty line (after a similar grossing up). This gave an overall poverty line of 1,160,842. Individual households lived in regions with different prices, so their expenditure per capita was first deflated, and then compared to this poverty line. The result was an estimated headcount poverty rate in Vietnam of 58% (World Bank 1999).

To compare poverty over time, this poverty line was updated to 1998. The cost of each item in the poverty-line diet of 1993 was recomputed using the prices of 1998 (as taken from the

[4] Pradhan et al. (2000) favor an interactive procedure: pick a reference population that is relatively poor and compute their cost of Calories; now recompute the poverty line; take as the new reference population those households close to this poverty line and re-calculate the cost of Calories; compute the poverty line again; and so on, until the poverty line stabilizes.

price questionnaire component of the VLSS, mainly); non-food expenditure was inflated using data from the General Statistical Office's price index. This yielded a poverty line of 1,793,903, and an associated poverty rate of 37%. The details are summarized in Table 3.6.

Table 3.6: Poverty Lines and Headcount Measures of Poverty, Vietnam

	Poverty line (m dong/capita/yr)	Headcount poverty rate
Poverty overall		
1993	1,160 ($109)	58
1998	1,790 ($135)	37
Food poverty		
1993	750 ($70)	25
1998	1,287 ($97)	15

Note: The food poverty rate excludes any provision for non-food items; it sets the poverty line at z^F.
Sources: Vietnam Living Standards Surveys of 1992-93 and 1997-98.

There is no wholly satisfactory way to measure the non-food component of the poverty line, and the procedures followed tend to be somewhat ad hoc. We saw above that for Vietnam, researchers essentially used the (slightly adjusted) level of non-food spending by households that were in the middle expenditure quintile in 1993. The poverty lines developed for South Korea (KIHASA 2000) measure the cost of food plus the cost of housing that meets the official minimum apartment size plus the cost of non-food items as measured by average spending by households in the poorest two-fifths of the income distribution.

Is there a better way to proceed? Probably not. Even the theory calls for compromise. Consider the food expenditure function shown in figure 3.6. Generally, $b = f(y)$, where b is food purchases and y is total expenditure. Following Ravallion (1998), let b^F be the cost of buying 2,100 Calories. Then an upper poverty line might be given by

(3.13) $$f^{-1}(b^F) = z^F$$

which measures the income level at which the household would buy 2,100 Calories of food; this is essentially the poverty line used in Vietnam. The non-food component is given by A (in figure 3.6).

A lower poverty line might be given by

(3.14) $$z_L^F = b^F$$

which measures the expenditure level at which the household could just buy enough food, but would not have any money left over to buy anything else; in Vietnam this is referred to as the food poverty line. But even in this case, households will typically buy non-food items, as shown by C in Figure 3.2. Ravallion

suggests that one might want to compromise, and measure non-food at the mid-point between these two extremes, giving *B*. In each case, the poverty line would be given by

(3.15) $z = b^F + 0$ *(or A or B')*

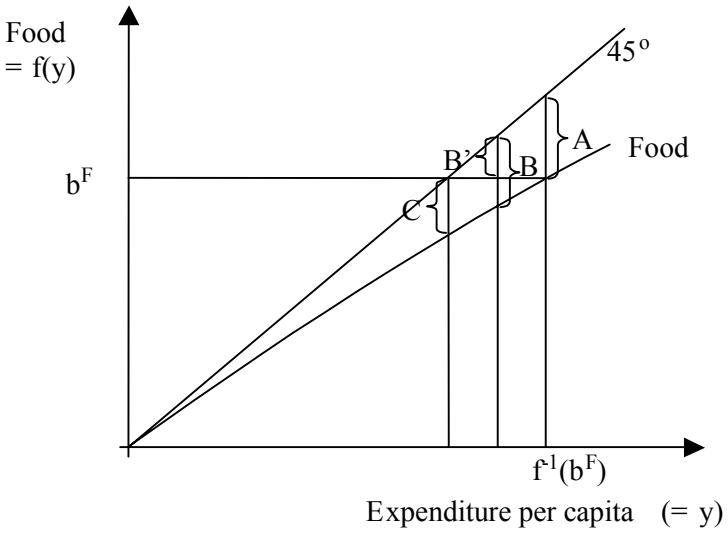

Figure 3.2 **Food Expenditure Function**.

9.		The non-food component of the poverty line, under the Cost of Basic Needs approach, may be obtained as
	A	The cost of basic housing and services.
	B	Non-food consumption of a household with just enough income to buy 2,100 Calories of food per capita per day along with other necessary goods and services.
	C	Non-food consumption of a household with just enough income to buy 2,100 Calories of food per capita per day.
	D	All of the above.

As one might expect when there is potential disagreement about the best approach to take, practice varies widely from one analyst to the next. Table 3.7 summarizes the approaches used to measure poverty in Africa, based on World Bank *Poverty Assessments* undertaken up to 1998. Based on a list of forty cases of poverty measurement compiled by Hanmer et al. (1999), 23 measured relative poverty; most of these set the poverty line as a share of mean income or expenditure (11 cases) or identified the poor using some percentage (e.g. 20%, 25%) of the income or expenditure distribution. The remaining 17 cases used an absolute measure of poverty, with most of them beginning with a calorie requirement (12 cases), sometimes adding a non-food component (5 cases). In a further five cases the analysts specified a basked of goods (including food) that was intended to measure the cost of basic needs but did not begin by identifying a calorie requirement. The heterogeneity of these measures makes it

difficult to compare poverty across countries, although in most country poverty assessments this is of secondary importance.

Table 3.7: Typology of poverty lines in World Bank Poverty Assessments for Africa		
Absolute (17 cases)	Calorie requirement (12)	Calories only Calorie cost/food share (1) Calories + basket of goods (5)
	Basket of goods (including food) (5)	
Relative (23 cases)	Relative to income base	Multiple of wage
		Share of mean income or expenditure (11)
	Specified percentage of income distribution (11)	
Source: Hanmer et al. (1999).		

3.3.2 Food Energy Intake method

The basic needs approach outlined above requires information on the prices of the goods that the poor consume. When price data are not available, a number of researchers have used an alternative method to construct the poverty line – the *food energy intake method*. As before, the goal here is to find the level of consumption expenditure (or income) that allows the household to obtain enough food to meet its energy requirements. Note that consumption will include non-food as well as food items; even underfed households typically consume some clothing and shelter, which means that at the margin these "basic needs" must be as valuable as additional food.

The basic idea is captured in figure 3.3, which shows a *calorie income function*; as income (or expenditure) rises, food energy intake also rises, although typically more slowly. Given some level of just-adequate food energy intake k, one may use this curve to determine the poverty-line level of expenditure, z. Formally, the function shows

(3.9) $$k = f(y)$$

So, given monotonicity,

(3.10) $$y = f^{-1}(k),$$

or, given a minimum adequate level of calorie k_{min}, we have

(3.11) $$z = f^{-1}(k_{\min})$$

where z is the poverty line. This approach is parsimonious in that it does not require any information about the prices of goods consumed.

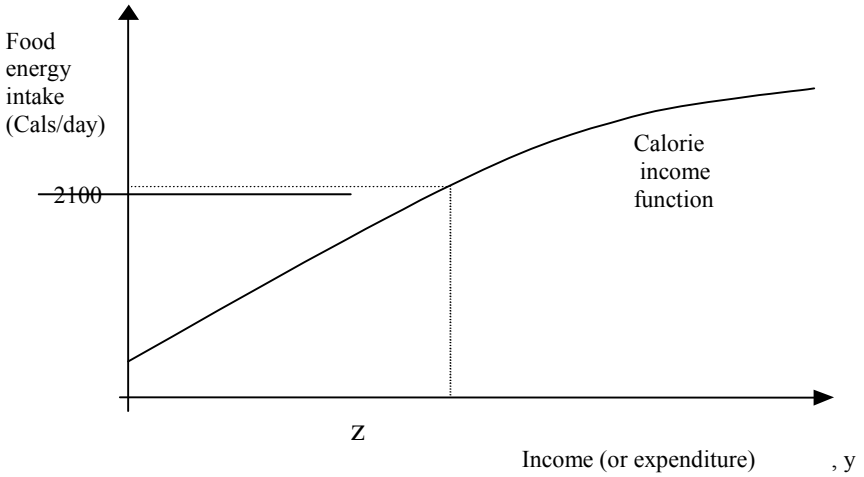

Figure 3.3: Calorie income function

First one needs to determine the amount of food that is adequate. Vietnam pegs this level at 2,100 Calories per person per day, in line with FAO recommendations, but it is recognized that individuals may need more or less food than this – clearly the needs of young children, growing teenagers, manual workers, pregnant women, or sedentary office workers may differ quite markedly; physical stature also plays a role. Not all countries have set the same cut-off point, as table 3.8 shows:

Table 3.8: Per Capita Daily Calorie Intake Used in Poverty Line Construction	
Vietnam	2,100
Indonesia	2,100
Philippines	2,000
Thailand	1,978
China	2,150

A variant of this approach was used to measure poverty in Vietnam, using data from the Vietnam Living Standards Survey of 1993. Separate food expenditure lines were estimated for urban and rural areas in each of seven provinces; the cost of obtaining 2,100 Calories of food per person per day was then computed, and the associated poverty lines – one for each rural and urban area in each province. This gave a headcount index of 55% (Dollar et al. 1995).

10.	Is the following statement true, false or uncertain? The Food Energy Intake approach sets the poverty line at the level of expenditure at which the household buys just enough Calories (e.g. 2,100 Calories per capita per day).		
	True	False	Uncertain

Unfortunately, the Food Energy Intake method is seriously flawed, and should not be used unless the alternatives are infeasible. Ravallion and Bidani (1994) computed headcount poverty measures for Indonesia using the SUSENAS data for 1990, both for the Cost of Basic Needs and the Food Energy Intake methods. Their results are shown in table 3.9. The most striking finding is that the poverty rates measured by the two approaches differ sharply! Ravallion and Bidani also computed poverty rates using these two measures, for each of the main regions of Indonesia, and found almost no correlation between the two measures.

Table 3.9: Headcount Measures of Poverty in Indonesia, 1990

	Cost of Basic Needs Method		Food Energy Intake Method
	Food	Food + non-food	
Indonesia overall	7.9	19.6	15.1
Urban	2.8	10.7	16.8
Rural	10.2	23.6	14.3
Source: Ravallion and Bidani 1994.			

Why is the Food Energy Intake method potentially unreliable? The weaknesses of the method were pointed out in an important article by Ravallion and Bidani (1994); in the next few paragraphs we summarize their approach and findings. The method also failed in a recent analysis of data from Vietnam, for slightly different reasons, also summarized below.

3.3.2.1 The urban-rural problem

The problem begins when one recognizes that food energy – typically shown on the Calorie income function – depends on other factors as well as income. The other influences include the tastes of the household (e.g. urban tastes in food may differ from rural tastes); the level of activity of household members; the relative prices of different foods, and of food to non-food items; and the presence of publicly-provided goods.

Figure 3.4 shows hypothetical (but plausible) calorie income functions for urban and rural households. Rural households can obtain food more cheaply, both because food is typically less

expensive in rural areas and also because they are more willing to consume foodstuffs that are cheaper per calorie (such as cassava rather than rice); urban consumers are more likely to buy higher quality foodstuffs, which raises the cost per calorie. It follows that the calorie income function for rural households will typically be higher than that of urban households. The implication is that for a given level of food energy intake, the poverty line in the rural area will be lower than in the urban area, as figure 3 makes clear. To the extent that this reflects differences in the cost of living, it is not a problem to have two poverty lines of this kind.

The key finding of Ravallion and Bidani (1994), based on 1990 data from the SUSENAS household survey in Indonesia, was that the urban poverty line (20,614 rupiah per person per month) was *much* higher than the rural one (13,295 Rp./person/month). This gap far exceeded the difference in the cost of living between urban and rural areas. Using these poverty lines, Ravallion and Bidani (1994) found that that poverty in Indonesia appeared to be higher in the urban than in the rural areas (Table 3.10), a completely implausible result. The point is also illustrated in figure 3.5, which shows the cumulative distribution of consumption per capita, for rural and urban areas, and marks the poverty lines and headcount poverty rates.

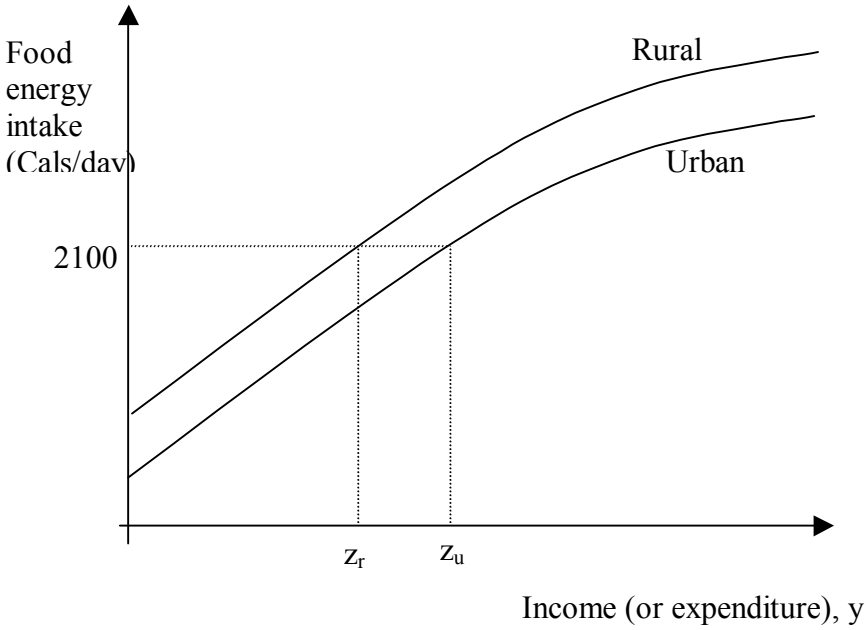

Figure 3.4. Calorie income functions for urban and rural Indonesia

Table 3.10: Poverty Lines in Indonesia using Food Energy Intake Method, 1990			
	Indonesia overall	Urban areas	Rural areas
P_0 (%)	15.1	16.8	14.3
P_1 (%)	2.42	3.23	1.06
P_2 (× 100)	0.66	0.94	0.53
Source: Ravallion and Bidani 1994			

.

11.	Ravallion and Bidani found, using the Food Energy Intake method, that the urban poverty line in Indonesia exceeded the rural poverty line by more than a simple comparison of living costs would lead one to expect, because	
	A	Urban households eat more.
	B	Urban households eat better quality food.
	C	Urban food prices are much higher than rural food prices.
	D	Urban housing costs more than rural housing.

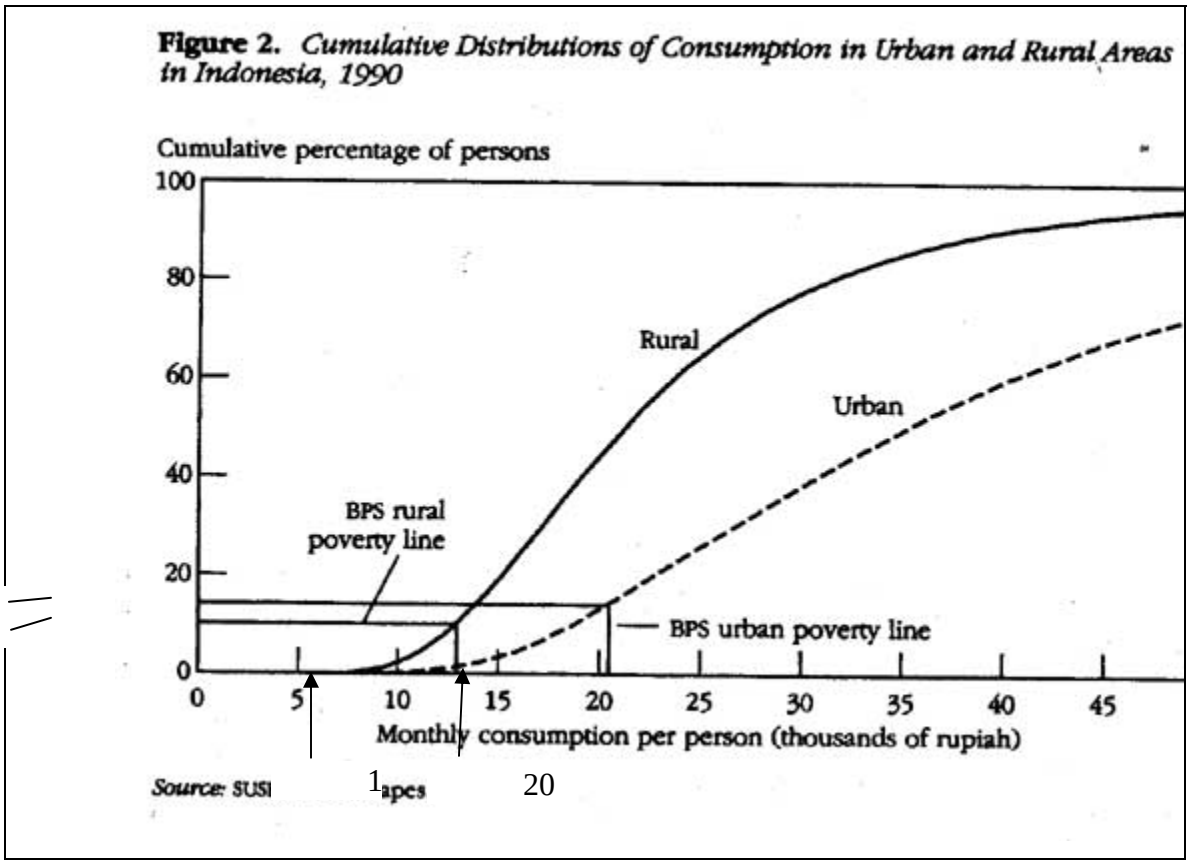

Figure 3.5 Cumulative distribution functions for consumption, Indonesia, 1990
(*Source*: Ravallion and Bidani 1994.)

3.3.2.2 The relative price problem

When researchers tried to apply the Food Energy Intake approach to data from the Vietnam Living Standards Survey of 1998, the method failed. As with the 1993 data, the idea was to compute food expenditure functions, find the cost of 2,100 Calories of food, and then find the related level of expenditure per capita, which would then serve as a poverty line. After undertaking this exercise, researchers found a higher level of poverty in 1998 than in 1993, an implausible result in an economy whose real GDP grew by 9% annually between 1993 and 1998, and where there was a widespread sense that the benefits of this growth had spread widely.

What went wrong? Figure 3.6 shows the situation. The food expenditure function shifted down between 1993 and 1998; in other words, for a given (real) income, households in 1998 would buy less food than in 1993. The main reason was that the price of food rose by 70% between 1993 and 1998, while the price of non-food items rose by just 25%; in other words, food became *relatively* much more expensive. As a result consumers shifted away from food to non-food consumption. This meant that the poverty line rose from z_{93} to z_{98} (see figure 5), a jump that turned out to be implausibly large.

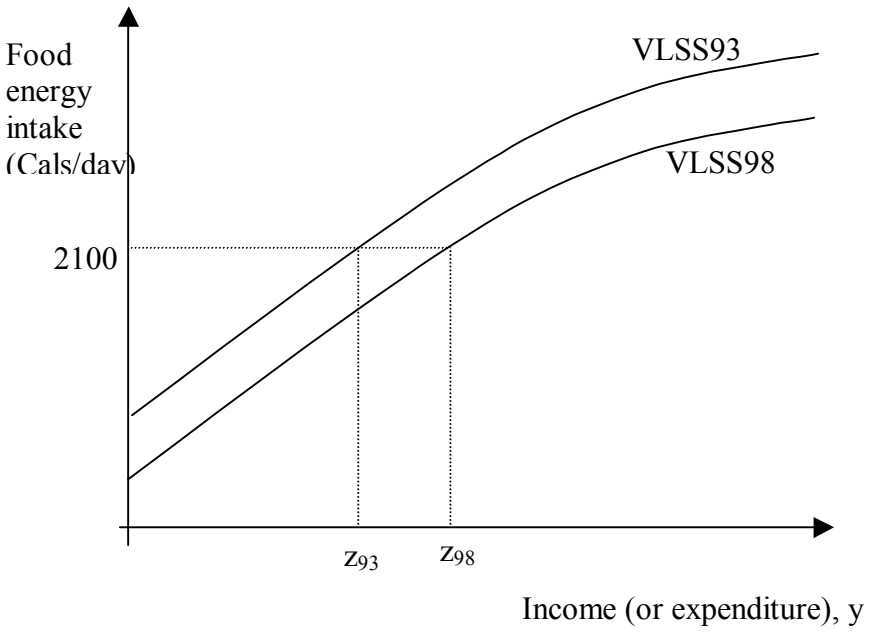

Figure 3.6 The determination of poverty lines for Vietnam, 1993 and 1998

This is a serious indictment of the Food Energy Intake method. But it should also be clear that every measure of poverty can be faulted, because it rests in part on arbitrary assumptions. In measuring poverty, there is no single truth.

12.	Is the following statement true or false? The Food Energy Intake method showed that the real poverty line in Vietnam rose rapidly between 1993 and 1998, because of inflation.	
	True	False

Solution B: Subjective poverty lines

We could measure poverty by asking people to define a poverty line, and using this to measure the extent of poverty. For instance, in a survey one might ask

"What income level do you personally consider to be absolutely minimal? That is to say that with less you could not make ends meet."

The answers will vary from person to person (and by size of household), but they could be plotted, and a line fitted through them, to get a subjective poverty line such as z^* in figure 3.7. It may also be possible to get adequate results by asking "do you consider your current consumption to be adequate to make ends meet?"

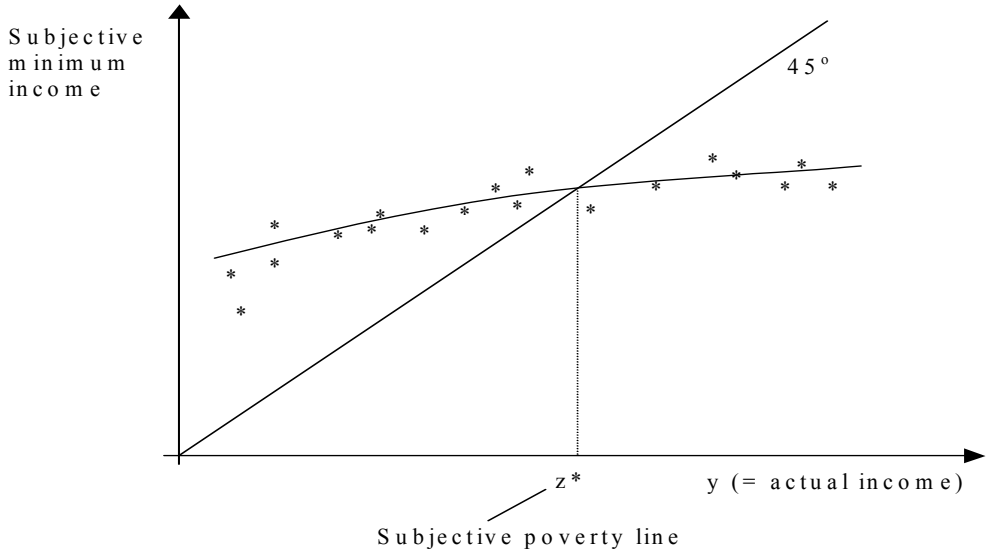

Figure 3.7 **Estimating a subjective poverty line.**

Mahar Mangahas has amassed extensive information on subjective poverty in the Philippines as part of the social weather stations project. Collected biannually since 1985, and quarterly since 1992, the surveys poll about 1,200 households. Each household is shown a card with a line running across it; below the line is marked *poor* ("mahirap") and above the line *non-poor*, and each household is asked to mark on the card where it fits. Separately, households are also asked to define a poverty line. Figure 3.8 reproduces a graph that tracks the evolution of this poverty rate from 1983 to 2003. Here are the comments of Mahar Mangahas that accompany the graph (Mangahas 2003, p.2):

> "The proportion of household heads rating their families as mahirap or poor was 62% in September, compared to a very low 53% in June, implying a return, roughly speaking, to conditions in November 2003 when Self-Rated Poverty was 61% [see Figure 3.8]. … Among poor households, the national median poverty threshold, or home expense budget needed in order not to feel poor, as of September 2003, is a modest P8,000 per month (P14,000 in Metro Manila, P8,000 elsewhere in Luzon, P5,000 in the Visayas, and P5,000 in Mindanao). This means that these home budgets are sufficient to satisfy one-half of the poor."[5]

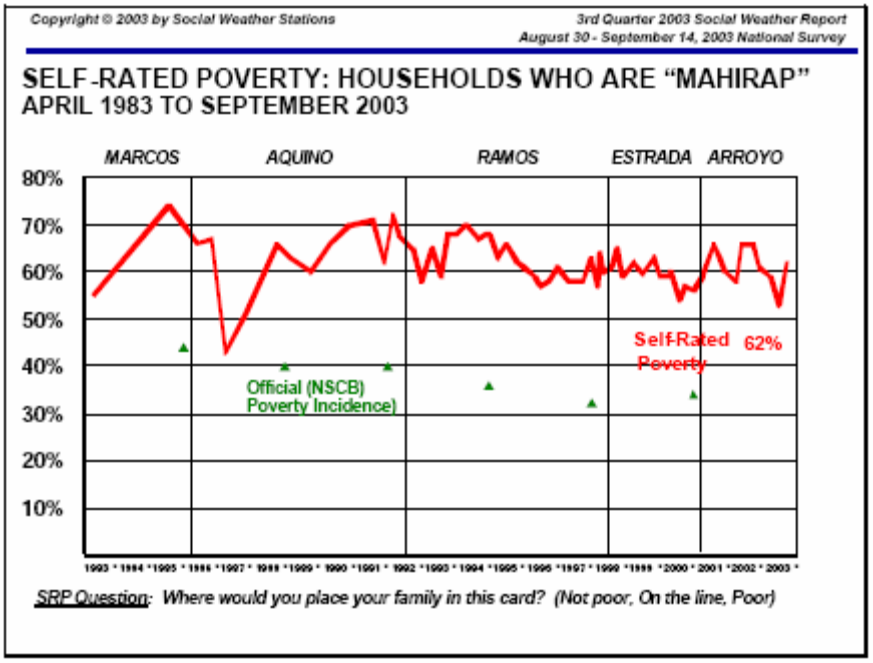

Gaurav Datt of the World Bank has analyzed the Filipino data in some detail. Here are some of his more interesting findings:

- *Self-rated poverty lines are high.* In 1997 the median poverty line was about 10,000 pesos per month for a "typical" household; this compares with the government's "basic needs" poverty line, which at

[5] The exchange rate in mid-September 2003 was P54.75/USD.

that time stood at 4,495 pesos/month. The implication is that self-rated poverty rates are high – 60% of all households in 1997, compared to 25% using the basic needs line.

- *The self-rated poverty line has risen rapidly over time,* by about 60-70% between 1985 and 1997. One consequence is that there is no trend in self-rated poverty over time. Another implication is that even when there is an economic slowdown, as occurred in 1997-98, the self-rated poverty rate hardly changes: it rose from 59% in 1996-97 to 61% in 1998.
- Perhaps surprisingly, the *self-rated poverty line given by poor households is only slightly lower* than that for non-poor households, and in fact the difference is not statistically significant. One might have expected poor households to have a less generous measure of the poverty line.
- There is a *clear urban/rural difference in perceptions* of the poverty line, with urban households setting a (money) poverty line at about twice the level of rural households, giving:

(3.16) $$z^{u}_{self-rated} \approx 2\, z^{r}_{self-rated}$$

The cost of living is certainly higher in urban areas, but by a factor of 1.2-1.5 rather than by a factor of 2. Thus the urban self-rated poverty line is, in real terms, higher than its rural counterpart. Why?

- One possibility is that there is more inequality in the urban areas, and that this raises expectations.
- Another plausible explanation is that households in urban areas may have more exposure to the media, and may have been affected more thoroughly by consumerism.
- A third explanation is that urban households may be more attuned to political processes, and their estimates of the poverty line may include an element of strategic behavior – trying to influence policy makers.

Self-rated measures of poverty are rarely collected. If the Filipino experience is at all representative, it is clear that they cannot usefully supplant the more traditional "objective" measures of poverty.

13.		Based on experience in the Philippines, which of the following statements is *not* true?
	A	Subjective poverty lines are not absolute over time.
	B	Self-rated poverty lines show high poverty rates.
	C	The rich report markedly higher poverty lines than the poor.
	D	Urban households set poverty lines higher than rural households, by more than the price differential between urban and rural areas would imply.

References

Joseph Dalaker and Mary Naifeh. 1998. *Poverty in the United States: 1997*, U.S. Department of Commerce, Washington DC.

Datt, Gaurav. [on social weather station project]

Dollar et al. 1995 [Vietham]

Gibson. 1999.

Haughton. 2000.

Kakwani. 2000?

KIHASA. 2000.

Mangahas. 2003.

Pradhan et al. 2000. [Indonesia]

Prescott and Pradhan. 1997.

Martin Ravallion. 1998. *Poverty Lines in Theory and Practice*, Living Standards Measurement Study Working Paper No. 133, World Bank.

Martin Ravallion and Benu Bidani. 1994. "How Robust Is a Poverty Profile?", *The World Bank Economic Review,* 8(1): 75-102.

Ravallion, Datt and van de Walle. 1991.

Sen. 1983.

World Bank. 1999. *Entering the 21st Century: World Development Report 1999/2000.* Washington DC.

World Bank. 2000. [Absolute and relative poverty rates.]

Vietnam: General Statiistical Office. 2000. *Viet Nam Living Standards Survey 1997-1998*, Statistical Publishing House, Hanoi.

World Bank. 2000. *World Development Report 2000-2001.* Washington DC.

Lars Osberg and Kuan Xu. 1999. *Poverty Intensity: How Well Do Canadian Provinces Compare?" Canadian Public Policy* (XXV:2).

Arsenio Balisacan. 1992. "xxx", *Journal of Philippine Development*, 19(1). [On estimating equivalence scales.]

Martin Ravallion. 1996. "How well can method substitute for data? Five experiments in poverty analysis," *The World Bank Research Observer*, 11(2), 199-221, August.

Vu Tuan Anh. 1997. "Criteria for Rural Poverty Monitoring," *Vietnam's Socio-Economic Development*, Spring.

Christina Paxson. 1993. "Consumption and Income Seasonality in Thailand," *Journal of Political Economy*, 101(1): 39-72.

Harold Alderman and Christina Paxson. 1994. "Do the Poor Insure? A Synthesis of the Literature on Risk and Consumption in Developing Countries," in Edmar Bacha (ed.), *Economics in a changing world: Proceedings off the Tenth World Congress of the International Economic Association*, Moscow, 48-78.

Chapter 4. Measures of Poverty

Summary

Assume that information is available on a welfare measure such as income per capita, and a poverty line, for each household or individual. This chapter explains how one may then construct summary measures of the extent of poverty.

The *headcount index* (P_0) measures the proportion of the population that is poor. It is popular because it is easy to understand and measure. But it does not indicate how poor the poor are.

The *poverty gap index* (P_1) measures the extent to which individuals fall below the poverty line (the poverty gaps) as a proportion of the poverty line. The sum of these poverty gaps gives the minimum cost of eliminating poverty, if transfers were perfectly targeted. The measure does not reflect changes in inequality among the poor.

The squared *poverty gap ("poverty severity") index* (P_2) averages the squares of the poverty gaps relative to the poverty line. It is one of the Foster-Greer-Thorbecke (FGT) class of poverty measures that may be written as

$$P_\alpha = \frac{1}{N}\sum_{i=1}^{N}\left(\frac{G_i}{z}\right)^\alpha$$

where N is the size of the sample, z is the poverty line, G_i is the poverty gap and α is a parameter; when α is larger the index puts more weight on the position of the poorest.

The *Sen-Shorrocks-Thon index* is defined as

$$P_{SST} = P_0 P_1^P (1 + \hat{G}^P)$$

where P_0 is the headcount index, P_1^P is the poverty gap index for the poor only, and G^P is the Gini index for the poverty gaps for the whole population. This measure allows one to decompose poverty into three components and to ask: Are there more poor? Are the poor poorer? And is there higher inequality among the poor?

Other measures of poverty are available. The *time taken to exit* measures the average time it would take for a poor person to get out of poverty, given an assumption about the economic growth rate; it may be obtained as the Watts Index divided by the growth rate of income (or expenditure) of the poor.

Learning Objectives

After completing the module on *Measures of Poverty*, you should be able to:

c. Describe and explain the headcount index, indicate why it is popular, and explain why it is an imperfect measure of poverty.
d. Describe and compute the poverty gap and poverty severity indexes, and evaluate their adequacy as measures of poverty.
e. Explain and evaluate the FGT (Foster-Greer-Thorbecke) family of poverty measures.
f. Compute the Sen and Sen-Shorrocks-Thon indexes of poverty, and show how the latter may be decomposed to identify the sources of changes in poverty.
g. Compute the Watts index and the related Time-Taken-To-Exit measure.
h. Argue that there is no single best measure of poverty.

Given information on a welfare measure such as per capita consumption, and a poverty line, then the only remaining problem is deciding on an appropriate summary measure of aggregate poverty. There are a number of aggregate measures of poverty that can be computed. The formulas presented here are all based on the assumption that the survey represents a simple random sample of the population, which makes them relatively easy to understand. Where the sampling is more complex – the typical situation in practice – weighting is needed, and the relevant formulas and associated programming are somewhat more difficult, but can be handled fairly easily by most major statistical packages such as Stata and SPSS.

4.1 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_0. Formally,

(4.1)
$$P_0 = \frac{N_p}{N}$$

where N_p is the number of poor and N is the total population (or sample). If 60 people are poor in a survey that samples 300 people, then $P_0 = 60/300 = 0.2 = 20\%$. For reasons that will be clearer below, it is often helpful to rewrite (4.1) as

(4.2)
$$P_0 = \frac{1}{N}\sum_{i=1}^{N} I(y_i < z),$$

Here, *I(.)* is an indicator function that takes on a value of 1 if the bracketed expression is true, and 0 otherwise. So if expenditure (y_i) is less than the poverty line (z), then *I(.)* equals to 1 and the household would be counted as poor. N_p is the total number of the poor.

The greatest virtues of the headcount index are that it is simple to construct and easy to understand. These are important qualities. However the measure has at least three weaknesses:

First, the headcount index does not take the intensity of poverty into account. Consider the following two income distributions:

Headcount Poverty Rates in A and B, assuming poverty line of 125					
	Expenditure for each individual in country				Headcount poverty rate (P_0)
Expenditure in country A	100	100	150	150	50%
Expenditure in country B	124	124	150	150	50%

Clearly there is greater poverty in country A, but the headcount index does not capture this. As a welfare function, the headcount index is unsatisfactory in that it violates the transfer principle – an idea

first formulated by Dalton (1920) that states that transfers from a richer to a poorer person should improve the measure of welfare. Here if a somewhat poor household were to give to a very poor household, the headcount index would be unchanged, even though it is reasonable to suppose that poverty overall has lessened.

Some argue that if it is to meaningful, the headcount index should imply that there is a "jump" or discontinuity in the distribution of welfare at about the poverty line, so it makes sense to speak of the poor and the non-poor. In practice, such a jump is not found (Ravallion 1996, p.1330).

Second, the head-count index does not indicate how poor the poor are, and hence does not change if people below the poverty line become poorer. Moreover, the easiest way to reduce the headcount index is to target benefits to people just below the poverty line, because they are the ones who are cheapest to move across the line. But by most normative standards, people just below the poverty line are the least deserving of the poor.

Third, the poverty estimates should be calculated for individuals and not households. If 20% of households are poor, it may be that 25% of the population is poor (if poor households are large) or 15% are poor (if poor households are small); the only relevant figures for policy analysis are those for individuals.

But survey data are almost always related to households, so in order to measure poverty at the individual level we must make a critical assumption that all members of a given household enjoy the same level of well-being. This assumption may not hold in many situations. For example, some elderly members of a household, or girls, may be much poorer than other members of the same household. In reality, not all consumption is evenly shared across household members.

4.2 Poverty gap index

A moderately popular measure of poverty is the *poverty gap index*, which adds up the extent to which individuals on average fall below the poverty line, and expresses it as a percentage of the poverty line. More specifically, define the poverty gap (G_i) as the poverty line (z) less actual income (y_i) for poor individuals; the gap is considered to be zero for everyone else. Using the index function, we have

(4.3) $$G_i = (z - y_i).I(y_i < z).$$

Then the poverty gap index (P_1) may be written as

(4.4) $$P_1 = \frac{1}{N}\sum_{i=1}^{N}\frac{G_i}{z}.$$

This table shows how the poverty gap is computed, divided by the poverty line, and averaged to give P_1, the poverty gap index.

Calculating the Poverty Gap Index, assuming poverty line of 125					
	Expenditure for each individual in country				Poverty Gap Index (P_1)
Expenditure in country C	100	110	150	160	
Poverty gap	25	15	0	0	
G_i/z	0.20	0.12	0	0	0.08 [= 0.32/4]

This measure is the mean proportionate poverty gap in the population (where the non-poor have zero poverty gap). Some people find it helpful to think of this measure as the cost of eliminating poverty (relative to the poverty line), because it shows how much would have to be transferred to the poor to bring their incomes or expenditures up to the poverty line (as a proportion of the poverty line). The minimum cost of eliminating poverty using targeted transfers is simply the sum of all the poverty gaps in a population; every gap is filled up to the poverty line. However this interpretation is only reasonable if the transfers could be made perfectly efficiently, for instance with lump sum transfers, which is implausible. Clearly this assumes that the policymaker has a lot of information; one should not be surprised to find that a very "pro-poor" government would need to spend far more than this in the name of poverty reduction.

At the other extreme, one can consider the maximum cost of eliminating poverty, assuming that the policymaker knows nothing about who is poor and who is not. From the form of the index, it can be seen that the ratio of the minimum cost of eliminating poverty with perfect targeting (i.e.G_i) to the maximum cost with no targeting (i.e. z, which would involve providing everyone with enough to ensure they are not below the poverty line) is simply the poverty gap index. Thus this measure is an indicator of the potential saving to the poverty alleviation budget from targeting: the smaller is the poverty gap index, the greater the potential economies for a poverty alleviation budget from identifying the characteristics of the poor – using survey or other information – so as to target benefits and programs.

The poverty gap measure has the virtue that it does not imply that there is a discontinuity ("jump") at the poverty line. To see this, consider the following example:

Poverty Gap Poverty Rates in A and B, assuming poverty line of 125						
	Expenditure for each individual in country				Poverty gap rate (P_1)	Headcount index (P_0)
Expenditure in country A	99	101	150	150	0.10	50%
Expenditure in country B	79	121	150	150	0.10	50%

For both of these countries, the poverty gap rate is 0.10, but most people would argue that country B has more serious poverty because it has an extremely poor member. Alternatively, one could think of the distribution in A as being generated from that in B by transferring 20 from the poorest person to the next poorest person – hardly an improvement in most people's eyes, yet one that has no effect on the poverty gap rate!

4.3 Squared poverty gap ("poverty severity") index

To construct a measure of poverty that takes into account inequality among the poor, some researchers use the squared poverty gap index. This is simply a weighted sum of poverty gaps (as a proportion of the poverty line), where the weights are the proportionate poverty gaps themselves; a poverty gap of (say) 10% of the poverty line is given a weight of 10% while one of 50% is given a weight of 50%; this is in contrast with the poverty gap index, where they are weighted equally. Hence, by squaring the poverty gap index, the measure implicitly puts more weight on observations that fall well below the poverty line. Formally:

(4.5)
$$P_2 = \frac{1}{N}\sum_{i=1}^{N}\left(\frac{G_i}{z}\right)^2.$$

This table shows how the poverty gap is computed, divided by the poverty line, squared, and averaged to give P_2, the squared poverty gap index.

Calculating the Squared Poverty Gap Index, assuming poverty line of 125					
	Expenditure for each individual in country				Squared Poverty Gap Index (P_2)
Expenditure in country C	100	110	150	160	
Poverty gap	25	15	0	0	
G_i/z	0.20	0.12	0	0	
$(G_n/z)^2$	0.04	0.0144	0	0	0.0136 [= 0.0544/4]

The measure lacks intuitive appeal, and because it is not easy to interpret it is not used very widely. It may be thought of as one of a family of measures proposed by Foster, Greer and Thorbecke (1984), which may be written, quite generally, as

(4.6)

$$P_\alpha = \frac{1}{N}\sum_{i=1}^{N}\left(\frac{G_i}{z}\right)^\alpha, \qquad (\alpha \geq 0)$$

where α is a measure of the sensitivity of the index to poverty and the poverty line is z, the value of expenditure per capita for the i-th person's household is x_i, and the poverty gap for individual i is $G_i=z-x_i$ (with $G_i=0$ when $x_i>z$) When parameter $\alpha=0$, P_0 is simply the head-count index. When $\alpha=1$, the index is the poverty gap index P_1, and when α is set equal to 2, P_2 is the poverty severity index. For all $\alpha > 0$, the measure is strictly decreasing in the living standard of the poor (the lower your standard of living, the poorer you are deemed to be). Furthermore, for $\alpha > 1$ it also has the property that the increase in measured poverty due to a fall in one's standard of living will be deemed greater the poorer one is. The measure is then said to be "strictly convex" in incomes (and "weakly convex" for $\alpha=1$). Another convenient feature of the FGT class of poverty measures is that they can be disaggregated for population sub-groups and the contribution of each sub-group to national poverty can be calculated.

Although the Foster, Greer and Thorbecke measure provides an elegant unifying framework for measures of poverty, it leaves unanswered the question of what is the best value of α. Moreover some of these measures also lack emotional appeal.

The measures of poverty depth and poverty severity provide complementary information on the incidence of poverty. It might be the case that some groups have a high poverty incidence but low poverty gap (when numerous members are just below the poverty line), while other groups have a low poverty incidence but a high poverty gap for those who are poor (when relatively few members are below the poverty line but with extremely low levels of consumption). Table 4.1 provides an example from Madagascar. According to the headcount measure (P_0), unskilled workers show the third highest poverty rate, while the group is in the fifth rank according to the poverty severity index (P_2). Compared to herders, they have a higher risk of being in poverty, but their poverty tends to be less severe. The types of interventions needed to help the two groups are therefore likely to be different.

Table 4.1: Poverty Indices By sub-groups, Madagascar, 1994

	Head count: % P_0	Rank	Poverty gap: % P_1	Rank	Poverty severity: × 100 P_2	Rank
Small farmers	81.6	1	41.0	1	24.6	1
Large farmers	77.0	2	34.6	2	19.0	2
Unskilled workers	**62.7**	**3**	**25.5**	**4**	**14.0**	**5**
Herders/fishermen	51.4	4	27.9	3	16.1	3
Retirees/handicapped	50.6	5	23.6	5	14.1	4

Source: Coudouel, Hentschel and Wodon (2001)

4.4 Sen Index.

Sen (1976) has proposed an index that sought to combine the effects of the number of poor, the depth of their poverty, and the distribution of poverty within the group. The index is given by

(4.7) $$P_s = P_0(1-(1-G^P)\frac{\mu^P}{z}),$$

where P_0 is the headcount index, μ^P is the mean income (or expenditure) of the poor, and G^P is the Gini coefficient of inequality among the poor. The Gini coefficient ranges from 0 (perfect equality) to 1 (perfect inequality), and is discussed in chapter 5 in the context of measuring inequality. The Sen Index can also be written as the average of the headcount and poverty gap measures, weighted by the Gini coefficient of the poor, giving:

(4.8) $$P_s = P_0 G^P + P_1(1-G^P)$$

It can be shown (Osberg and Xu 2002) that the Sen Index may also be written as

(4.9) $$P_S = P_0 P_1^P (1+G^{PP}),$$

where G^{PP} is the Gini coefficient of the poverty gap ratios of only the poor and P_1^P is the poverty gap index *calculated over poor individuals only*.

The Sen index has been widely discussed, and has the virtue of taking the income distribution among the poor into account. However the index is almost never used outside of the academic literature, perhaps because it is lacks the intuitive appeal of some of the simpler measures of poverty, but also because it "cannot be used to decompose poverty into contributions from different subgroups" (Deaton, 1997, p.147).

4.5 The Sen-Shorrocks-Thon index.

The Sen index has been modified by others, and perhaps the most compelling version is the Sen-Shorrocks-Thon (SST) index, defined as

(4.10) $$P_{SST} = P_0 P_1^P (1+\hat{G}^P),$$

which is the product of the headcount index, the poverty gap index (*applied to the poor only*), and a term with the Gini coefficient of the poverty gap ratios (i.e. of the G_n's) *for the whole population*. This Gini coefficient typically is close to 1, indicating great inequality in the incidence of poverty gaps.

Example. In 1996, 12.4% of the population of Quebec province (Canada) was in poverty. The poverty gap index, *applied to the poor only*, stood at 0.272. And the Gini coefficient of the poverty gap ratios was 0.924. Thus the Sen-Shorrocks-Thon index was 0.065 (=0.124 × 0.272 × (1+0.924)).

Application. Osberg and Xu (1999) use the SST index to compare poverty across the 10 Canadian provinces for 1984, 1989, 1994, 1995 and 1995, as well as to put the degree of Canadian provincial poverty into an international context. A number of graphs from their study are reproduced below. Figure 4.1 provides an international comparison, using the SST index, and shows that the US is an outlier with its relatively high poverty rate (as measured by the SST). A comparison of the US and Canada over time (figure 4.2) shows that while poverty was similar in the two countries a generation ago, it is now clearly higher in the US than in Canada. Figure 4.3 provides information on some Canadian provinces: Newfoundland was the poorest in 1984, but by 1996 had become much less of an outlier.

One strength of the SST index is that it can help give a good sense of the sources of change in poverty over time. This is because the index may be decomposed into

(4.11) $$\Delta \ln P_{SST} = \Delta \ln P_0 + \Delta \ln P_1^P + \Delta \ln(1 + \hat{G}^P),$$

which may be interpreted as, % change in SST index = % change in headcount index + % change in poverty gap index (among poor) + % change in (1+Gini coefficient of poverty gaps).

In plain English, this allows us to decompose poverty into three aspects: are there more poor? are the poor poorer? and is there higher inequality among the poor?

FIGURE 1
1990s – Country Rankings by SST Poverty Index – Poverty Line Based on Half the Median Equivalent Income

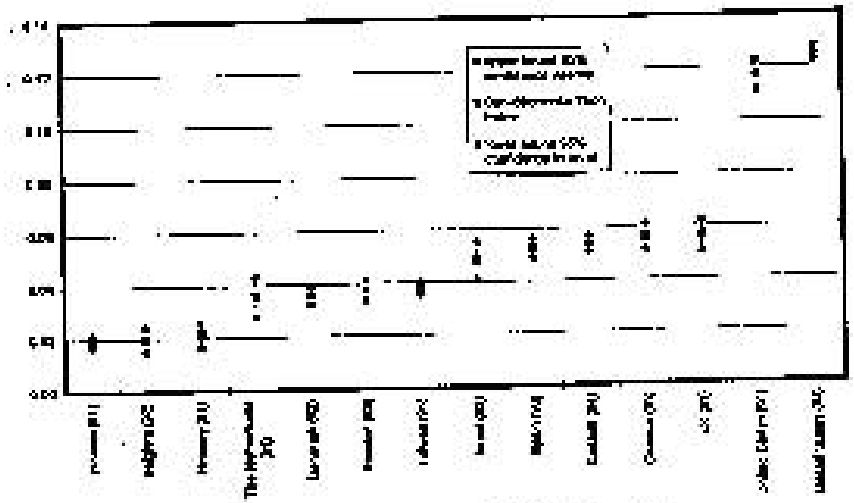

Note: [95% confidence interval = mean +/-2 standard deviations] of 500 bootstraps.
Source: Luxembourg Income Study, Osberg and Xu (1997).

FIGURE 2
Canada-United States Comparison – The SST Index from 1971 to 1994 Poverty Line Based on Half the Median Equivalent Income

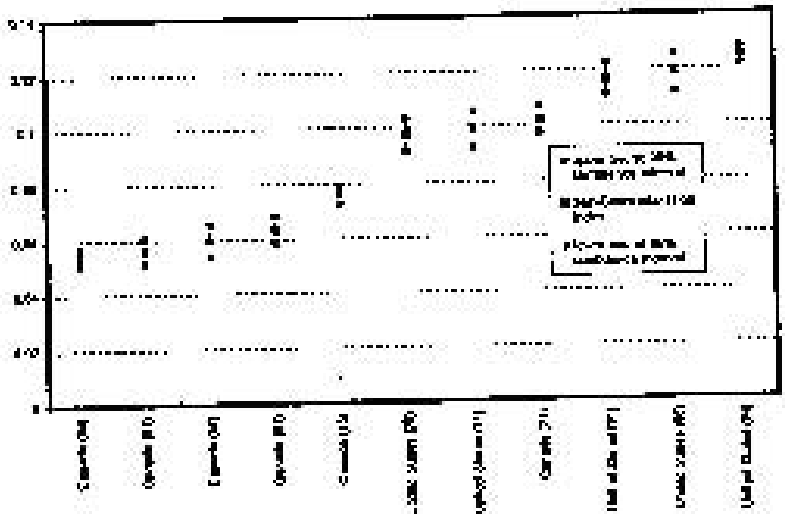

Note: [95% confidence interval = mean +/-2 standard deviations] of 500 bootstraps.
Source: Luxembourg Income Study, Osberg and Xu (1997)

FIGURE 3
Sen-Shorrocks-Thon Index of Poverty Intensity – 1/2 Median After-Tax Income, 1984, 1989, 1996

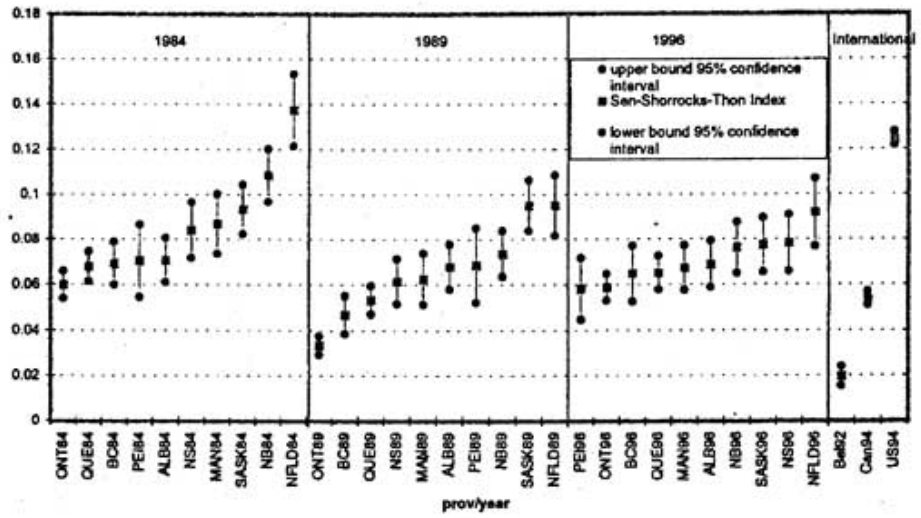

Note: [95% confidence interval = mean +/-2 standard deviations] of 300 bootstraps.
Source: Authors' calculations, Survey of Consumer Finances, various years. Luxembourg Income Study, Osberg and Xu

Example. The information in table 4.2 comes from Osberg and Xu, and traces the evolution of poverty in the Canadian province of Newfoundland between 1984 and 1996. It is clear that most of the change in the poverty rate over time was due to variations in the number of people in poverty (P_1), rather than in the size of the poverty gap per poor person (P_1^P) or the distribution of poverty among the poor (G^P).

Table 4.2: Decomposition of poverty, and changes in poverty, in Newfoundland, 1984-1996

	SST index	P_0	P^P_1	$1+G^P$	ΔlnSST index	ΔLnP_0	ΔlnP^P_1	$\Delta\ln(1+G^P)$
1984	.137	.245	.304	1.844				
1989	.095	.169	.296	1.897	-.370*	-.372*	-.027	.028
1994	.105	.184	.304	1.884	.104	.086	.026	-.007
1995	.125	.212	.316	1.864	.168	.141	.038	-.010
1996	.092	.164	.294	1.897	-.307	-.254	-.071	.018

Notes: * denotes statistically significant at the 95% level. Poverty line is half of median equivalent income, using the "OECD scale" (i.e. equivalent income = $1 + 0.7(N_{adults}-1)+0.5(N_{children})$.
Source: Osberg and Xu, 1999.

Note that the values of the Sen-Shorrocks-Thon index provided by Osberg and Xu do not give just a single point estimate for each province; they also provide a confidence interval. Because the SST

index is complex, it is not possible to compute these confidence intervals analytically. Instead, they are computed artificially using *bootstrapping*. The basic idea behind the bootstrap is straightforward and clever. Suppose we have a survey sample of 2,000 households. Now pick a sample of 2,000 from this sample *with replacement* – i.e. pick a household, then put it back into the sample, pick another household, put it back into the sample, and so on, until you have picked 2,000 households. Some households will be chosen more than once, but that's fine. Now compute the SST index using this artificial sample. Then repeat the process many times; Osberg and Xu use 300 repetitions. The result is a distribution of values of the SST, from which it is easy to find (say) the 95% confidence interval. Sample Stata code to generate confidence intervals for the SST index is given in the Appendix, in the exercises associated with Chapter Five.

4.6 The Watts Index

The first distribution-sensitive poverty measure was proposed in 1968 by Watts (see Zheng 1993), and in its discrete version takes the form:

(4.13)
$$W = \frac{1}{N}\sum_{i=1}^{q}[\ln(z) - \ln(y_i)]$$

where the N individuals in the population are indexed in ascending order of income (or expenditure), and the sum is taken over the q individuals whose income (or expenditure) y_i falls below the poverty line z.

This table shows how the Watts index is computed, by dividing the poverty line by income, taking logs, and finding the average over the poor. The Watts index is attractive in that it satisfies all the theoretical properties that one would want in a poverty index, and is increasingly used by researchers in generating such measures as the poverty incidence curve (see chapter xxx).[6] However, it is not a particularly intuitive measure, and so is rarely seen in practical field work.

Calculating the Watts Index, assuming poverty line of 125					
	Expenditure for each individual in country				Watts Index
Case 1 (poor)					
Expenditure in country C	100	110	150	160	
z/y_i	1.25	1.14	0.83	0.78	
log (z/y_i)	0.223	0.128	-0.182	-0.247	**0.351**
Case 2 (less poor)					
Expenditure in country C	110	120	150	160	

[6] Ravallion and Chen (2001) argue that three axions are essential to any good measure of poverty. Under the focus axiom the measure should not vary if the income of the non-poor varies); under the monotonicity axiom, any income gain for the poor should reduce poverty; and under the transfer axiom, inequality-reducing transfers among the poor should reduce poverty. The Watts index satisfies these three axioms, but the headcount (P_0) and poverty severity (P_1) measures do not.

z/y_i	1.14	1.04	0.83	0.78	
log (z/y_i)	0.128	0.041	-0.182	-0.247	**0.169**
Case 3 (deeper poverty)					
Expenditure in country C	90	120	150	160	
z/yi	1.25	1.10	0.83	0.78	
log (z/yi)	0.329	0.041	-0.182	-0.247	**0.369**

4.7 **Time taken to exit**

Most poverty profiles for Cambodia, and indeed for most countries, rely on the three basic classes of Foster Greer Thorbecke poverty statistics discussed above. But when thinking about poverty reduction strategies, it may be useful to show how long it would take, at different potential economic growth rates, for the average poor person to exit poverty. A poverty statistic with this property is derived by Morduch (1998); the statistic is decomposable by population sub-groups and is also sensitive to how expenditure (or income) is distributed among the poor. For the *j*th person below the poverty line, the expected time to exit poverty (i.e., to reach the poverty line), if consumption per capita grows at positive rate *g* per year is:

(4.12)
$$t_g^j \approx \frac{\ln(z) - \ln(x_j)}{g} = \frac{W}{g}.$$

In other words, the time take to exit is the same as the Watts index divided by the expected growth rate of income (or expenditure) of the poor.

What effect can economic growth have on the elimination of poverty? Figure 4.4 shows the average time it would take to raise the consumption level of a poor person in Cambodia to the poverty line, for various hypothetical growth rates. It is assumed that this growth rate is continuous, is in real terms, and is distributionally neutral *among* the poor. If the economic growth rate enjoyed by the poor were only one percent per year, it would take over 20 years for the average poor person to exit poverty. But at a growth rate of four percent per year it would take less than six years for the average poor person to exit poverty. Hence, economic growth that acts to raise the real consumption levels of the poor can have a powerful effect on the elimination of poverty.

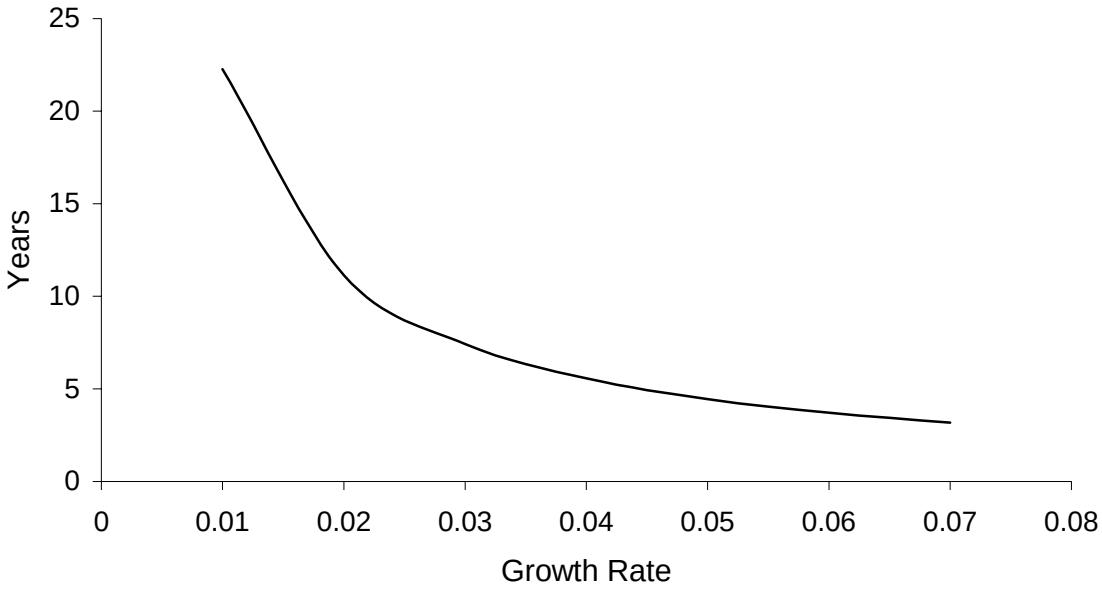

Figure 4.4: ***Average exit time from poverty***

Despite the potency of economic growth, it will generally take more than just growth to rapidly improve the lives of the very poor. The expected time to exit poverty for those people who are so poor that they are below the food poverty line in Cambodia – i.e. they cannot afford enough food, even if they were to devote all their consumption spending to food – is more than 15 years, even at a three percent continuous annual growth rate. Thus, targeted programs are needed to deliver benefits to the poor, for instance in the form of improvements in their human and physical assets or through interventions (e.g., infrastructure, markets) that improve the returns they get from those assets.

4.8 Other Measures

There are other additive poverty measures that are distribution sensitive. Following Atkinson (1987), one can characterize a general class of additive measures, encompassing W, the FGT (Foster, Greer and Thorbecke) class of measures, and some other measures (such as the second measure proposed by Clark, Hemming and Ulph, 1981), as taking the following form:

(4.14)
$$P = \frac{1}{N}\sum_{i=1}^{N} p(z, y_i)$$

where p(z, yi) is the individual poverty measure, taking the value zero for the non-poor (y_i>z) and some positive number for the poor, the value of which is a function of both the poverty line and the individual living standard, non-decreasing in the former and non-increasing in the latter.

Given the wide variety of aggregate measures of poverty that are available, which ones should one use? We turn to this question in chapter 5.

References

Vu Tuan Anh. 1997. "Criteria for Rural Poverty Monitoring," *Vietnam's Socio-Economic Development*, Spring.

Atkinson, Anthony. 1987. "On the measurement of poverty," *Econometrica*, 55: 749-764.

Clark, Hemming and Ulph. 1981.

Dalton, Hugh. 1920. "The measurement of the inequality of incomes," *Economic Journal*, 30: 384-361.

Deaton, Angus. 1997. *The Analysis of Household Surveys*, Johns Hopkins University Press, Baltimore MD.

Foster, James, J. Greer and Eric Thorbecke. 1984. "A class of decomposable poverty measures," *Econometrica*, 52(3): 761-765.

Morduch, Jonathan. 1998. "Poverty, economic growth, and average exit time," *Economics Letters*, 59: 385-390.

Lars Osberg and Kuan Xu. 1999. *Poverty Intensity: How Well Do Canadian Provinces Compare?" Canadian Public Policy/Analyse de Politiques*, XXV(2) :

Osberg and Xu. 2002.

Martin Ravallion. 1996. "How well can method substitute for data? Five experiments in poverty analysis," *The World Bank Research Observer*, 11(2), 199-221, August.

Ravallion, Martin and Shaohua Chen. 2001. *Measuring Pro-Poor Growth*, Development Research Group, World Bank.

Sen, Amartya. 1976. "Poverty: an ordinal approach to measurement," *Econometrica*, 46: 437-446.

Zheng, B. 1993. "An axiomatic characterization of the Watts poverty index," *Economics Letters*, 42(1): 81-86.

Chapter 5. Poverty Indices: Checking for Robustness

Summary

There are four main reasons why measures of poverty may not be robust.

Sampling error occurs because measures of poverty are based on sample data, which gives the true poverty rate only with some degree of uncertainty. It is good practice to report standard deviations and confidence intervals for poverty measures; this can be done by *bootstrapping*. Since household surveys tend to be relatively small, it is not possible to disaggregate the results to units smaller than relatively broad regions.

Measurement error occurs in all survey data; we know, for instance, that households under-report income and expenditure, which tends to overstate the degree of poverty. The effect can be large: in some cases a 5% understatement of consumption can translate into a 10% overstatement of the headcount poverty rate.

Poverty rates vary depending on the equivalence scale used, although the variation is typically fairly modest. Equivalence scales are not widely used, because of the difficulty of agreeing on an appropriate set of weights.

The choice of a poverty line, and associated poverty rate (e.g. headcount index, poverty gap index), is arbitrary. Sometimes, although not always, these choices matter. By comparing the cumulative distribution function of expenditure (or income) per capita – sometimes called the *poverty incidence curve* – between two situations, one may judge whether the choice of poverty line affects the conclusion about the change in poverty. If there is first order stochastic dominance, the choice of poverty line is not crucial; otherwise it is often possible to use higher order tests (e.g. second order stochastic dominance) to help reach a clear conclusion about whether poverty differs between two time periods (or regions or countries).

No study of poverty is complete without some discussion of the robustness of the findings.

Learning Objectives

After completing the module on *Poverty Indices: Checking for Robustness*, you should be able to:

i. Explain what is meant by robustness and why poverty measures might not be robust.
j. Describe sampling error, and argue the case for presenting standard deviations and confidence intervals along with poverty rates.
k. Explain what bootstrapping is and how it may be used to generate confidence intervals and sample standard deviations.
l. Enumerate the sources of measurement error.
m. Define the elasticity of the headcount index with respect to errors in mean expenditure (or income) per capita, and explain how this translates an understatement of expenditure (or income) into an overstatement of poverty.
n. Explain what an adult equivalence scale is and describe some common equivalence scales.
o. Explain why equivalence scales are not widely used in practice.
p. Define, and show how to graph, the poverty incidence curve and the poverty deficit curve.
q. Explain what is meant by first order stochastic dominance, and why it is useful when assessing how robust a poverty comparison is to the choice of poverty line or poverty measure.
r.

5.1 Introduction

Between 1998 and 2003, the poverty rate in Vietnam, as measured by the headcount index, fell from 37% to 29%. Good news indeed! But before celebrating, it is important to ask how robust this conclusion is.

The problem is that there is pervasive uncertainty about four possibly crucial aspects of a poverty comparison.

- First, there is ***sampling error,*** which arises because we are trying to measure the poverty rate for the country as a whole on the basis of information from a relatively modest number of households.
- Second, there is likely to be ***measurement error***, so that our measures of poverty are inherently imprecise.
- Third, there are unknown ***differences in needs*** between households at similar consumption levels, yet considerable arbitrariness in the way in which equivalence scales are used to address this problem.
- And fourth, there is ***uncertainty and arbitrariness about both the poverty line and the precise poverty measures*** used.

Given these problems, how robust are poverty comparisons? Would they change if we made alternative assumptions? These are the questions addressed in this chapter.

5.2 Sampling Error

Suppose I want to determine the average height of men in a country. It is unrealistic to measure everyone's height, so I pick a random sample of men and measure their heights. I find that the average height of men in the sample (the sample mean, $\bar{x}$) is 1.752 meters. For now, this is my best estimate of the average height of all men in the population. But if I were to take another random sample, I would almost certainly find a slightly different sample mean, 1.756 meters for instance. In other words, the sample mean is a random variable; to the extent that it only approximates the true mean height of men in the population, we have a degree of sampling error.

Let s be the standard deviation of heights in the sample (e.g. 0.2 meters). Then the estimated standard deviation of the sample mean is given by $s/\sqrt{n}$, where n is the number of observations in the sample (e.g. 0.01 meters, if $n = 400$). It follows that as the sample becomes larger, the standard deviation of the sample mean decreases – in other words, the estimate becomes more precise and sampling error decreases. Note that it takes a quadrupling of the sample size to halve the standard deviation of the

sample mean. Acceptably accurate measures of income and expenditure levels, and poverty rates, generally require a sample size of close to 1,000; thus a national survey of 6,000 households only allows the results to be disaggregated into about 6-8 regions.

Measures of poverty are more complex than a simple estimate of the height of men in a country, and the estimation of the standard deviation of the poverty measure is usually difficult to obtain analytically. The solution is to use the technique of *bootstrapping*. As explained in Chapter 4, one computes the poverty rate many times, by sampling with replacement from the survey data. From the distribution of these computed poverty rates it is possible to measure the standard deviation and create confidence intervals (such as the interval within which 95% of the estimated measures of poverty fall). Some sample Stata code for bootstrapping is given in the Appendix.

Whether the standard deviations and confidence intervals are computed analytically or using bootstrapping, the important point is that they should be computed and presented along with the estimated poverty rates, so that the reader has a clear sense of the magnitude of the sampling error. It is more honest, useful and accurate to say that "our best estimate of the headcount poverty rate is 32.2% and we are 95% confident that the true rate falls in the range 29.9% to 34.6%" than to say "the poverty rate is 32.2%".

5.3 Measurement error

Poverty measures can be quite sensitive to certain sorts of measurement error in the underlying parameters, and quite robust to others. For example, suppose that the observed welfare indicator contains an additive random error with zero mean; one gets the indicator's value right on average, but there is an error in any one observation. Suppose also (for the purpose of this example) that the poverty line is at the mode of the distribution. Then it can be shown that the expected value of the observed head-count index will be locally un-affected by changes in the degree of imprecision in the welfare measurements; on average one will predict the same index with a noisy indicator as a precise one, and the estimated headcount poverty rate will be unbiased. However, this need not hold for other poverty lines or for higher order poverty measures (such as the poverty gap measure); Ravallion (1988) gives a general statement of the necessary and sufficient conditions for greater variability in the welfare indicator to increase the expected value of a poverty measure defined on that indicator. Under plausible assumptions, greater imprecision in the welfare indicator will yield higher estimates of any smooth distribution-sensitive poverty measure, such as P_2, in effect overstating poverty as a result.

Now consider, instead, errors in the mean of the distribution on which the poverty measure is being estimated; suppose, for instance, that consumption is systematically underestimated, as is typical, due to incomplete recall, deliberate omissions, or the exclusion from the survey of poor but hard-to-reach groups such as the homeless. It can be demonstrated that the elasticity of the head-count index to errors in the mean, holding the income distribution (as measured by the Lorenz curve) constant, is simply the elasticity of the cumulative distribution function evaluated at the poverty line.[7] This may be estimated, and values around two are quite common, at least for developing countries. Thus a five percent under-estimation of mean consumption at all consumption levels could easily translate into a 10 percent over-estimation in the head-count index and, hence, the number of poor. Furthermore, amongst the Foster-Greer-Thorbecke class of measures, the elasticities tend to be higher, the higher the value of alpha (which is 0 for the headcount ratio and 1 for the poverty gap ratio).

In practice, inflation – defined as a substantial, sustained increase in the general level of prices – raises the nominal value both of expenditure and of the poverty line, and so has no net effect on the measures of poverty.

5.4 Equivalence scales

Poverty studies usually measure living standards using expenditure (or income) per capita. As discussed in Chapter Two, because needs vary among household members, and because there are economies of scale in consumption, poverty measures based on per capita welfare indicators may not be good estimates. An alternative is to base our poverty measures on expenditure (or income) *per adult equivalent*. If poverty estimates are not affected by the adult equivalence weights that we choose, it is safe to say that those poverty estimates are not biased as a consequence of the weighting procedure used.

Kathleen Short and her colleagues have investigated the sensitivity of the US 1997 poverty rate to a number of variations on the OECD equivalence scales (Short et al., *Experimental Poverty Measures*, June 1999). The results are summarized in the table below. While different definitions of adult equivalent do change the measured poverty rate in the United States, the most striking feature of this table is how small these differences are. In other words, the choice of adult equivalent scale may not matter very much.

[7] The elasticity of Y with respect to X is the percentage change of Y divided by the percentage change of X. The result is a number that is unit-free.

Table 5.1. Sensitivity of Headcount Poverty Rate(P_0) to Different Specifications of Adult Equivalence Scales, United States, 1999.	
Adult Equivalent Scale Used	Headcount poverty rate, %
N_a+N_c (i.e. use income per capita)	13.3
$(N_a + 0.7\ N_c)^{0.65}$	13.1
$(N_a + 0.7\ N_c)^{0.70}$	12.3
$(N_a + 0.7\ N_c)^{0.75}$	12.7
$(N_a + 0.7\ N_c)^{0.5}$	13.4
$(N_a + 0.7\ N_c)^{0.6}$	12.7
$(N_a + 0.85\ N_c)^{0.65}$	12.7
$1 + 0.4(N_a - 1) + 0.4$(child 1) $+ 0.3\ (N_c - 1)$ (Canadian scale)	13.8
One adult: 1. Two adults: 1.41. Single parents: $(N_a + 0.8 + 0.5\ N_c - 1)$. All other families: $(N_a + 0.5\ N_c)^{0.7}$	13.1
Source: Short et al., 1999. *Notes:* N_a = number of adults. N_c = number of children.	

Table 5.2 displays adult equivalence weights that have been used in India and in Taiwan. A study of the importance of weights computed the correlation coefficient between (unweighted) expenditure per capita and expenditure per adult equivalent using these weights, using data from Sri Lanka (1969-70), Taiwan (1974) and Peninsular Malaysia (1973). All the correlation coefficients were 0.96 or higher; they are shown in table 5.3 below.. Since the equivalence scales give similar results to those using expenditure per capita, the case for using adult equivalence scales (rather than the much simpler expenditure per capita) is not compelling. This helps explain why adult equivalence scales are not used more often in practical poverty analysis.

Table 5.2: Adult equivalents, India and Taiwan					
India: adult equivalents			Taiwan: adult equivalents		
Age	Male	Female	Age	Male	Female
0	.43	.43	**0-1**	.3	.3
1-3	.54	.54	**2-4**	.4	.4
4-6	.72	.72	**5-7**	.5	.5
7-9	.87	.87	**8-10**	.7	.7
10-12	1.03	.93	**11-14**	.8	.8
13-15	.97	.80	**15-20**	.9	.9
16-19	1.02	.75	**21+**	1.0	.9
20-39	1.00	.71			
40-49	.95	.68			
20-59	.90	.64			
60-69	.80	.51			
70+	.70	.50			

Table 5.3: Correlation coefficients, expenditure per capita with expenditure per adult equivalent			
Survey	Years	India weights	Taiwan weights
Sri Lanka	1969-70	0.99	0.96
Taiwan	1974	0.98	0.96
Peninsular Malaysia	1973	0.99	0.97

5.5 Choice of Poverty Line and Poverty Measure

The choice of a poverty line, and the associated poverty measure (e.g. the headcount index P_0 or the poverty gap index P_1) is essentially arbitrary. However, if the various measures of poverty, introduced in chapter 4, tell the same story, then it does not matter much which measure one chooses, since they are close substitutes for one another.

In practice, however, the differences between measures can be quite pronounced. Consider, for example, two policies:

- Policy A entails a small redistribution from people around the mode of the distribution, which is also where the poverty line happens to be located, to the poorest households. (This is actually a fair characterization of how a reduction in the prices of domestically produced food-staples would affect the distribution of welfare in some Asian countries).
- Policy B entails the opposite change - the poorest lose while those near the mode gain. (An increase in food-staple prices in the above example).

A moment's reflection will confirm that the head-count index (H) will prefer policy B; since changes in H depend solely on which direction people are crossing the poverty line (under policy A, people who were just on the poverty line would now fall below it since they transfer resources to the very poor, thus the total number of poor increases). However, a measure such as the Squared Poverty Gap Index will indicate the opposite ranking, since it will respond relatively more to the gains among the poorest than among the not-so-poor. In short, the choice of poverty measure often does matter.

The need to examine higher order poverty measures, such as the Poverty Gap Index and the Squared Poverty Gap Index, also depends on whether or not the poverty comparison in terms of the head-count index has considered more than one poverty line. If only one poverty line is used then it is especially important to check the higher-order measures, as the above example illustrates. But if poverty has fallen no matter what poverty line is used, and no matter what measure (e.g. P_0, P_1, P_2) is used, then a finding of an increase in poverty is robust.

One can formalize the examination of robustness (to the choice of poverty line and/or poverty measure) by drawing on the theory of stochastic dominance. Ravallion (1998) gives a straightforward exposition of the approach, oriented to the needs of the analyst trying to make a reasonably robust poverty comparison.

5.5.1 A single measure of standard of living

As mentioned before, it is always a good idea to consider one or more extra poverty lines. Let us take this a step further, and imagine the curve that is traced out as one plots the head-count index (P_0) on the vertical axis and the poverty line on the horizontal axis, allowing the latter to vary from zero to the maximum consumption. This is simply the cumulative distribution function, and may also be thought of as the *poverty incidence curve,* F(z). Each point on the curve gives the proportion of the population consuming less than the amount given on the horizontal axis, as in Figure 5.1. Thus the poverty incidence curve for year 1 in Figure 5.1 shows that if the poverty line is $600, the poverty rate is 24%, if the poverty line is $900, the poverty rate is 53%, and so on.

Figure 5.1: ***Poverty Incidence Curves for years 1 and 2***

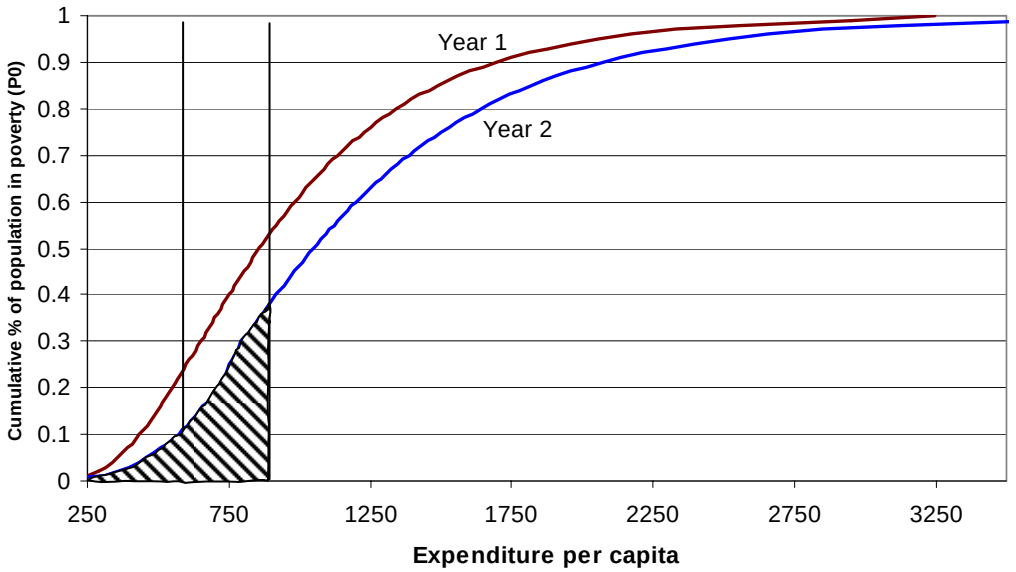

One can go further. If one calculates the area under the poverty incidence curve up to each point – the shaded area in Figure 5.2, for example – and then plots it against the poverty line, this traces out the *poverty deficit curve*, D(z), shown in figure 5.3. Each point on this curve is simply the total value of the poverty gap (or, equivalently, the poverty gap index times the poverty line z).

From Figure 5.1 we see that for every possible choice of poverty line, the poverty rate in year 2 is below that of year 1. Here there is *first order stochastic dominance*: the precise choice of poverty line is unimportant (at least up to the maximum conceivable poverty line z^{max}, which is the relevant range),

because no matter what poverty line is chosen, we still conclude that poverty fell between year 1 and year 2.

In Figure 5.1 the ranking was unambiguous. This is not so in Figure 5.2: using a poverty line of $600, poverty is higher in year 1; but with a poverty line of $1,000, poverty is higher in year 2 (as measured by the headcount index). In this case it is not at all clear whether poverty has risen or fallen.

Figure 5.2: ***Poverty Incidence Curves for years 1 and 2***

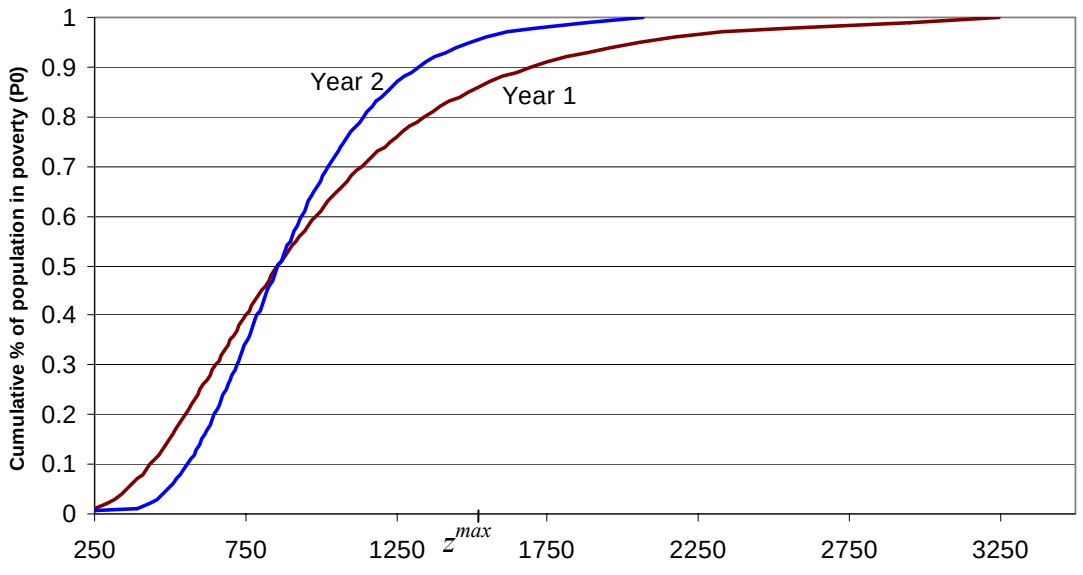

It is sometimes possible to resolve this issue by appealing to second (or higher!) order stochastic dominance. To do this, one must exclude the head-count index, and consider only additive measures that reflect the depth of poverty such as the poverty gap index (P_1) and the squared poverty gap index (P_2). A fall in poverty then requires that the *poverty deficit curve*, given by the area under the cumulative distribution, be nowhere lower for year 1 at all points up to the maximum poverty line, and at least somewhere higher. What happens above the z^{max} is not relevant, and in fact the overall probability need not be smaller under the lower poverty deficit curve, and could in even be higher; this case is illustrated in figure 5.3.

Figure 5.3: ***Poverty Deficit Curves for years 1 and 2***

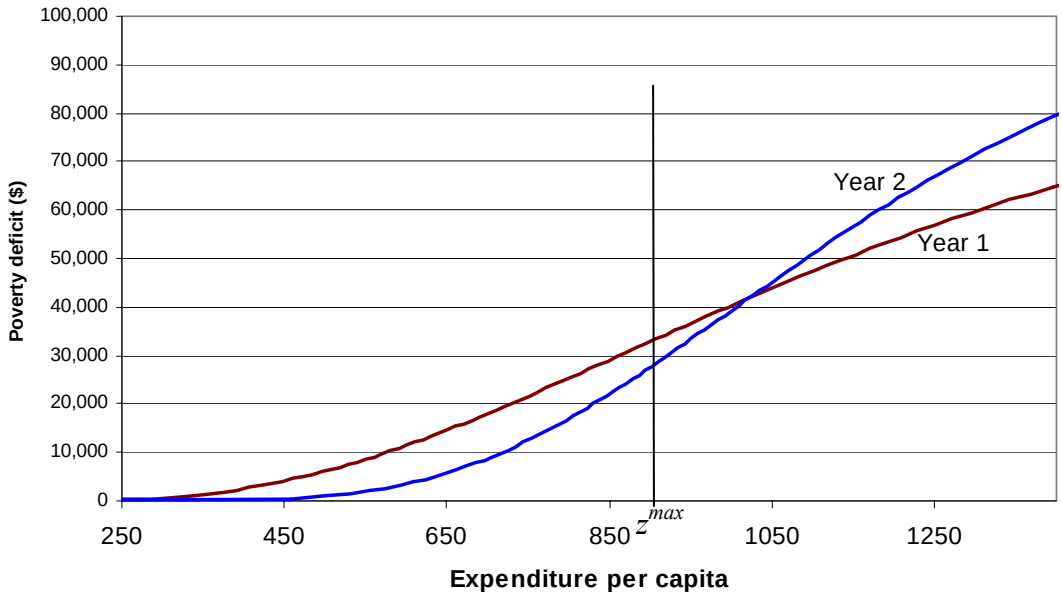

In this example, there is *second-order dominance in the relevant range*, and we may consider that poverty has indeed fallen. Intuitively, if the measure of poverty is the poverty gap index, this is equivalent to saying that the sum of the poverty gaps (i.e. the poverty deficit) is smaller in year 2 than in year 1, no matter what poverty line is used, provided it is below z^{max}.

To illustrate the two dominance tests, consider an initial state in which four persons have consumption in amounts (100, 110, 140, 150) in year 1 and these change to (110, 112, 128, 150) in year 2. Has poverty changed between year 1 and year 2?

To help answer this question, consider the numbers in Table 5.4. If the poverty line is 100, then a quarter of the population is poor in year 1 and none in year 2, and it would appear that poverty has fallen. But if the poverty line is set at 130, then the poverty rate was 0.5 in year 1 and actually rose to 0.75 in year 2. In other words, whether poverty (as measured by the headcount rate) has risen or fallen turns out to be sensitive to the choice of poverty line. Formally, the poverty incidence curves cross, so we do not have first order stochastic dominance.

Table 5.4. ***Poverty Incidence and Poverty Depth Curves***						
Poverty line (z) is:	**Poverty incidence curve, F(z)**			**Poverty depth curve, D(z)**		
	Year 1		Year 2	Year 1		Year 2

100	0.25	>	0.0	0.25	>	0.0
110	0.5	=	0.5	0.75	>	0.5
120	0.5	=	0.5	1.25	>	1.0
130	0.5	<	0.75	1.75	=	1.75
140	0.75	=	0.75	2.5	=	2.5
150	1.0	=	1.0	3.5	=	3.5
	NB: Poverty incidence curves cross			NB: Poverty deficit curves do not cross		
Note: Assumes a society of four people with consumption of (100, 110, 140, 150) in year 1 and (110, 112, 128, 150) in year 2.						

Now consider the poverty depth curve. If the poverty line is 120, then the value of the poverty depth curve is 1.25, obtained by cumulating the values of the poverty incidence curve (i.e. 0.25 + 0.5 + 0.5 in this case). Poverty depth was unambiguously higher (or at least not lower) in year 1 than in year 2 (see Table 5.4, final three columns), no matter what poverty line is chosen. In this case there is second order stochastic dominance, and we have a moderately robust finding that poverty has fallen. This makes some intuitive sense: between year 1 and year 2, average income did not change, but the incomes of the poorest members of society rose (albeit at the expense of one of the moderately-well off individuals). Such a redistribution would widely be considered to have reduced poverty. Similar comparisons can be made using other poverty measures, although this is not done often.

Technical aside: When two frequency distributions (e.g. poverty incidence curves) are quite close, we may also want to assess whether the difference between them is statistically significant, or potentially due simply to sampling error. For first-order dominance, this can be done using the Kolmogorov-Smirnov test, based on the largest vertical distance between the two cumulative frequency curves; most textbooks on statistics explain how to do this fairly straightforward test, and provide tabulations of critical values; the syntax in Stata may be found by typing `help ksmirnov` or `search kolmogorov`.

Box: First Order Stochastic Dominance, Formally

A more formal statement runs as follows: Consider two income distributions $\boldsymbol{y}_1$ and $\boldsymbol{y}_2$ with cumulative distribution functions (CDFs) $F(\boldsymbol{y}_1)$ and $F(\boldsymbol{y}_2)$. If $F(\boldsymbol{y}_1)$ lies nowhere above and at least somewhere below $F(\boldsymbol{y}_2)$ then distribution $\boldsymbol{y}_1$ displays first order stochastic dominance over distribution $\boldsymbol{y}_2$. $F(\boldsymbol{y}_1) \geq F(\boldsymbol{y}_2)$ for all y. Hence in distribution $\boldsymbol{y}_1$ there are no more individuals with income less than a given income level than in distribution $\boldsymbol{y}_2$, for all levels of income. We can express this in an alternative way using the inverse function $\boldsymbol{y}=F^{-1}(\boldsymbol{p})$ where $\boldsymbol{p}$ is the share of the population with income less than a given income level: first order dominance is attained if $F_1^{-1}(\boldsymbol{p}) \geq F_2^{-1}(\boldsymbol{p})$ for all $\boldsymbol{p}$. The inverse function $F^{-1}(\boldsymbol{p})$ simply plots incomes against cumulative population, usually using ranked income quintiles. First order stochastic dominance

of distribution y_1 over y_2 implies that any social welfare function that is increasing in income will record higher levels of welfare in distribution y_1 than in distribution y_2.

5.5.2 Robustness: More than One Dimension [Advanced]

Similar ideas can be applied in circumstances in which the poverty lines themselves vary across households or individuals in an unknown way – in other words, the poverty line itself is a random variable. For example, errors in measuring the standard of living may mean that we are not identifying the true poverty lines for different individuals. Unknown differences in "needs" at given consumption levels could also mean that the true poverty lines vary. There may be considerable, unknown, inter-individual variation in nutritional requirements. Errors in accounting for differences between households in their demographic composition or the prices they face may also entail some underlying variation in the appropriate poverty lines.

Poverty comparisons are clearly more difficult when the poverty line has an unknown distribution, but even then unambiguous conclusions may be possible if one is willing to make some assumptions. Provided that the *distribution* of poverty lines is the same for the two (or more) situations being compared and is independent of the distribution of living standards, first-order dominance of one distribution over another implies an unambiguous poverty ranking. This holds no matter what the underlying distribution of poverty lines may be.

Another case of interest is when one knows the distribution of "needs" (such as family size) as well as consumption, but one does not know precisely how these two variables interact to determine welfare. For two dimensions of welfare, such as aggregate consumption and family size, one can derive *bi-variate dominance tests*, which are more or less stringent depending on the assumptions one is willing to make about the way in which differences in "needs" interact with consumption in determining well-being; the precise tests depend on (amongst other things) whether the marginal social valuation of consumption is higher or lower in larger families. In the special case in which the marginal valuation of consumption is independent of family size, and the marginal distribution of size is fixed, the problem collapses back to the standard dominance tests above.

Let us suppose first that we know nothing about how needs interact with consumption in determining poverty. For additive poverty measures and a fixed distribution of the population across different needs, all of the above dominance tests can be applied separately to each of the groups identified

as having different needs. Thus one can test for first-order dominance (FOD) amongst (say) rural households, separately for urban households, or large families separately from small families. If we find that FOD holds for each group separately then we can conclude that FOD also holds for the aggregate, no matter what the difference in needs is between the groups. If FOD fails then, by restricting attention to measures of the depth and severity of poverty, one can then test for second-order dominance for each "needs" group separately, or third-order dominance if necessary.

More general tests can be devised for a large variety of situations, but many are difficult to explain non-mathematically and so are unlikely to be very convincing to policy makers.

References

Arsenio Balisacan. 1992. "xxx", *Journal of Philippine Development*, 19(1). [On estimating equivalence scales.]

Martin Ravallion. 1998. *Poverty Lines in Theory and Practice*, Living Standards Measurement Study Working Paper No. 133, World Bank.

Short, Kathleen et al. 1999. *Experimental Poverty Measures.*

Chapter 6. Inequality Measures

Summary

Inequality is a broader concept than poverty in that it is defined over the *entire* population, and does not only focus on the poor.

The simplest measurement of inequality sorts the population from poorest to richest and shows the percentage of expenditure (or income) attributable to each fifth (quintile) or tenth (decile) of the population. The poorest quintile typically accounts for 6-10% of all expenditure, the top quintile for 35-50%.

A popular measure of inequality is the Gini coefficient, which ranges from 0 (perfect equality) to 1 (perfect inequality), but is typically in the range of 0.3-0.5 for per capita expenditures. The Gini is derived from the Lorenz curve, which sorts the population from poorest to richest, and shows the cumulative proportion of the population on the horizontal axis and the cumulative proportion of expenditure (or income) on the vertical axis. While the Gini coefficient has many desirable properties – mean independence, population size independence, symmetry, and Pigou-Dalton Transfer sensitivity – it cannot easily be decomposed to show the sources of inequality.

The best-known entropy measures are Theil's T and Theil's L, both of which allow one to decompose inequality into the part that is due to inequality within areas (e.g. urban, rural) and the part that is due to differences between areas (e.g. the rural-urban income gap). Typically at least three-quarters of inequality in a country is due to within-group inequality, and the remaining quarter to between-group differences.

Atkinson's class of inequality measures is quite general, and is sometimes used. The decile dispersion ratio – defined as the expenditure (or income) of the richest decile divided by that of the poorest decile – is popular but a very crude measure of inequality.

It is often helpful to decompose inequality by occupational group, or by source of income, in order to identify policies that would help moderate inequality.

Learning Objectives

After completing the module on poverty lines, you should be able to:

21. Explain what inequality is, and how it differs from poverty.
22. Compute and display information on expenditure (or income) quintiles.
23. Draw and interpret a Lorenz curve.
24. Compute and explain the Gini coefficient of inquality.
25. Argue that the Gini Coefficient satisfies mean independence, population size independence, symmetry, and Pigou-Dalton Transfer sensitivity, but is not easily decomposable.
26. Compute and interpret generalized entropy measures, including Theil's T and Theil's L.
27. Compute and interpret Atkinson's inequality measure for different values of the weighting parameter ε.
28. Compute and criticize the decile dispersion ratio.
29. Decompose inequality using Theil's T in order to distinguish between-group from within-group components of inequality, for separate geographic areas, occupations, and income sources.

6.1 Definition of inequality

The main focus of this manual is on poverty, which looks at the situation of individuals or households who find themselves at the bottom of the income distribution; typically this requires information both about the mean level of (say) expenditure per capita as well as its distribution at the lower end. But sometimes we are more interested in measuring inequality than poverty per se, and for that reason we have included this relatively brief chapter on measuring inequality.

Inequality is a broader concept than poverty in that it is defined over the *entire* population, and not just for the population below a certain poverty line. Most inequality measures do not depend on the mean of the distribution, and this property of mean independence is considered to be a desirable property of an inequality measure. Of course, inequality measures are often calculated for distributions other than expenditure – for instance, for income, land, assets, tax payments, and many other continuous and cardinal variables.

The simplest way to measure inequality is by dividing the population into fifths (*quintiles*) from poorest to richest, and reporting the levels or proportions of income (or expenditure) that accrue to each level. Table 6.1 shows the level of expenditure per capita, in '000 dong per year, for Vietnam in 1993, based on data from the Vietnam Living Standards Survey. A fifth of the individuals (not households) included in the survey were allocated to each expenditure quintile. The figures show that 8.4% of all expenditures were made by the poorest fifth of households, and 41.4% by the top fifth. Quintile information is easy to understand, although sometimes one wants a summary measure rather than a whole table of figures.

Table 6.1: Breakdown of expenditure per capita by quintile, Vietnam 1993

	Expenditure quintiles					
	Lowest	Low-mid	Middle	Mid-upper	Upper	Overall
Per capita expenditure ('000 dong/year)	518	756	984	1,338	2,540	1,227
% of expenditure	8.4	12.3	16.0	21.8	41.4	100.0
Memo: Cumulative % of expenditure	8.4	20.7	36.7	58.5	100.0*	
Memo: Cumulative % of population	20.0	40.0	60.0	80.0	100.0	

Source: Vietnam Living Standards Survey 1993.
Note: * There is a slight rounding error here.

6.2 Commonly used summary measures of inequality

6.2.1 Gini coefficient of inequality

The most widely used single measure of inequality is the Gini coefficient. It is based on the Lorenz curve, a cumulative frequency curve that compares the distribution of a specific variable (e.g. income) with the uniform distribution that represents equality. To construct the Gini coefficient, graph the *cumulative* percentage of households (from poor to rich) on the horizontal axis and the *cumulative* percentage of expenditure (or income) on the vertical axis. The Lorenz curve shown in figure 1 is based on the Vietnamese data in Table 6.1. The diagonal line represents perfect equality. The Gini coefficient is defined as A/(A+B), where A and B are the areas shown on the graph. If A=0 the Gini coefficient becomes 0 which means perfect equality, whereas if B=0 the Gini coefficient becomes 1 which means complete inequality. In this example the Gini coefficient is about 0.35.

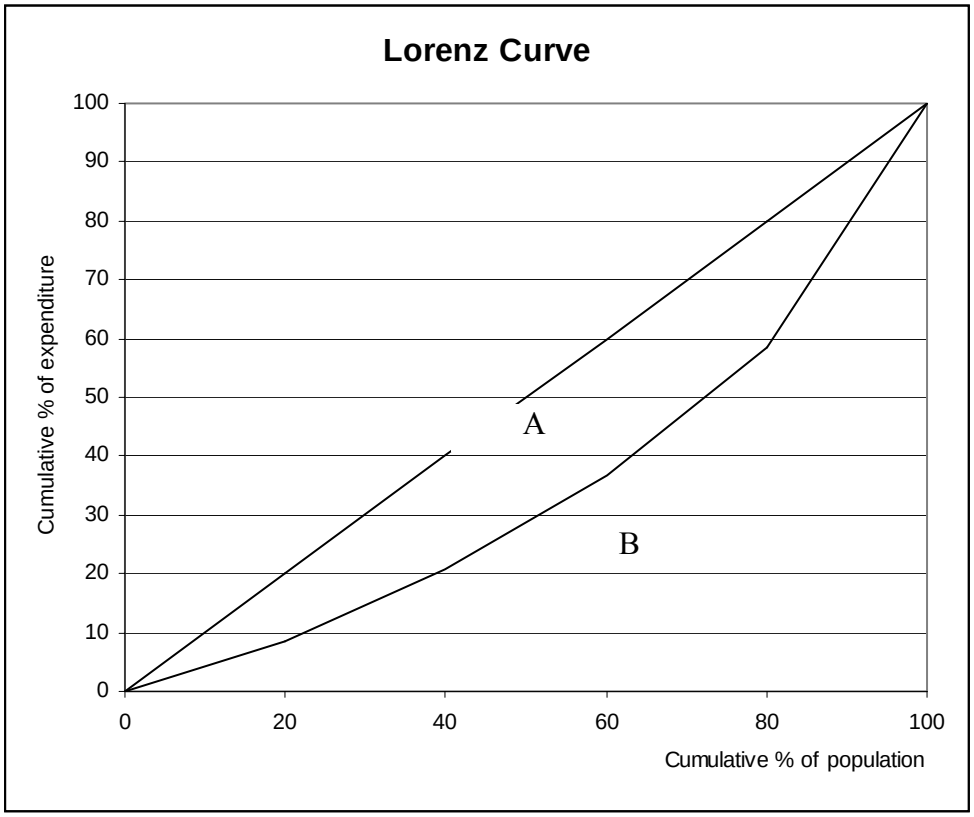

Figure 6.1. ***Lorenz Curve***

Formally, let x_i be a point on the X-axis, and y_i a point on the Y-axis. Then

(6.1) $$Gini = 1 - \sum_{i=1}^{N} (x_i - x_{i-1})(y_i + y_{i-1}).$$

When there are N equal intervals on the X-axis this simplifies to

(6.2) $$Gini = 1 - \frac{1}{N} \sum_{i=1}^{N} (y_i + y_{i-1}).$$

For users of Stata, there is a `gini` command that may be downloaded and used directly (see Appendix 3). This command also has the advantage that it allows one to use weights, which are not incorporated into the two equations shown above.

The Gini coefficient is not entirely satisfactory. To see this, consider the criteria that make a good measure of income inequality, namely:

- *Mean independence.* This means that if all incomes were doubled, the measure would not change. The Gini satisfies this.
- *Population size independence.* If the population were to change, the measure of inequality should not change, ceteris paribus. The Gini satisfies this too.
- *Symmetry.* If you and I swap incomes, there should be no change in the measure of inequality. The Gini satisfies this.
- *Pigou-Dalton Transfer sensitivity.* Under this criterion, the transfer of income from rich to poor reduces measured inequality. The Gini satisfies this too.

It is also desirable to have

- *Decomposability.* This means that inequality may be broken down by population groups or income sources or in other dimensions. The Gini index is not easily decomposable or additive across groups. That is, the total Gini of society is not equal to the sum of the Gini coefficients of its subgroups.
- *Statistical testability.* One should be able to test for the significance of changes in the index over time. This is less of a problem than it used to be because confidence intervals can typically be generated using bootstrap techniques.

6.2.2 Generalized Entropy measures

There are a number of measures of inequality that satisfy all six criteria. Among the most widely used are the Theil indexes and the mean log deviation measure. Both belong to the family of generalized entropy inequality measures. The general formula is given by:

(6.3)
$$GE(\alpha) = \frac{1}{\alpha(\alpha-1)}\left[\frac{1}{N}\sum_{i=1}^{N}\left(\frac{y_i}{\bar{y}}\right)^{\alpha} - 1\right]$$

where $\bar{y}$ is the mean income (or expenditure per capita). The values of GE measures vary between 0 and ∞, with zero representing an equal distribution and higher value representing a higher level of inequality. The parameter α in the GE class represents the weight given to distances between incomes at different parts of the income distribution, and can take any real value. For lower values of α, GE is more sensitive to changes in the lower tail of the distribution, and for higher values GE is more sensitive to changes that affect the upper tail. The commonest values of α used are 0,1 and 2. GE(1) is Theil's T index, which may be written as

(6.4)
$$GE(1) = \frac{1}{N}\sum_{i=1}^{N}\frac{y_i}{\bar{y}}\ln\left(\frac{y_i}{\bar{y}}\right)$$

GE(0), also known as Theil's L, and sometimes referred to as the mean log deviation measure, is given by:

(6.5)
$$GE(0) = \frac{1}{N}\sum_{i=1}^{N}\ln\left(\frac{\bar{y}}{y_i}\right)$$

Once again, users of Stata do not need to program the computation of such measures from scratch; the `GE` command, explained in Appendix 3, allows one to get these measures, even when weights need to be used with the data.

6.2.3 Atkinson's inequality measures

Atkinson has proposed another class of inequality measures that are used from time to time. This class also has a weighting parameter ε (which measures aversion to inequality) and some of its theoretical properties are similar to those of the extended Gini index. The Atkinson class, which may be computed in Stata using the `Atkinson` command, is defined as:

(6.6)

$$A_\varepsilon = 1 - \left[\frac{1}{N} \sum_{i=1}^{N} \left(\frac{y_i}{\bar{y}} \right)^{1-\varepsilon} \right]^{1/(1-\varepsilon)}, \quad \varepsilon \neq 1$$

$$= 1 - \frac{\prod_{i=1}^{N} (y_i^{(1/N)})}{\bar{y}}, \qquad \varepsilon = 1.$$

Table 6.2 sets out in some detail the computations involved in the computation of the Generalized Entropy and Atkinson measures of inequality. The first row of numbers gives the incomes of the ten individuals who live in a country, in regions 1 and 2. The mean income is 33. To compute Theil's T, one first computes $y_i/ybar$, where *ybar* is the mean income level; then compute ln($y_i/ybar$), take the product, add up the row, and divide by the number of people. Similar procedures yield other generalized entropy measures, and also the Atkinson measures.

Table 6.2: Computing Measures of Inequality

		Region 1					Region 2				
Incomes (=yi)		10	15	20	25	40	20	30	35	45	90
Mean income (ybar)	33.00										
yi/ybar		0.30	0.45	0.61	0.76	1.21	0.61	0.91	1.06	1.36	2.73
ln(yi/ybar)		-0.52	-0.34	-0.22	-0.12	0.08	-0.22	-0.04	0.03	0.13	0.44
Product		-0.16	-0.16	-0.13	-0.09	0.10	-0.13	-0.04	0.03	0.18	1.19
GE(1): Theil's T	**0.080**										
ln(ybar/yi)		0.52	0.34	0.22	0.12	-0.08	0.22	0.04	-0.03	-0.13	-0.44
GE(0): Theil's L	**0.078**										
(yi/ybar)^2		0.09	0.21	0.37	0.57	1.47	0.37	0.83	1.12	1.86	7.44
GE(2)	**0.666**										
(yi/ybar)^.5		0.55	0.67	0.78	0.87	1.10	0.78	0.95	1.03	1.17	1.65
Atkinson, e=0.5	**0.087**										
(yi)^(1/n)		1.26	1.31	1.35	1.38	1.45	1.35	1.41	1.43	1.46	1.57
Atkinson, e=1	**0.164**										
(yi/ybar)^(-1)		3.30	2.20	1.65	1.32	0.83	1.65	1.10	0.94	0.73	0.37
Atkinson, e=2	**0.290**										

Here are some examples of different measures of inequality (Dollar and Glewwe 1999, p.40):

Table 6.3: Expenditure inequality in selected less developed countries			
Country	Gini coefficient	Theil T	Theil L
Côte d'Ivoire, 1985-86	0.435	0.353	0.325
Ghana, 1987-88	0.347	0.214	0.205
Jamaica, 1989	n/a	0.349	0.320
Peru, 1985-86	0.430	0.353	0.319
Vietnam, 1992-93	0.344	0.200	0.169
Source: Reported in Dollar and Glewwe (1999), p.40.			

6.2.4 Decile dispersion ratio

A simple, and widely-used, measure is the decile dispersion ratio, which presents the ratio of the average consumption of income of the richest 10 percent of the population divided by the average income of the bottom 10 percent. This ratio can also be calculated for other percentiles (for instance, dividing the average consumption of the richest 5 percent – the 95th percentile – by that of the poorest 5 percent – the 5th percentile).

The decile ratio is readily interpretable, by expressing the income of the top 10% (the "rich") as a multiple of that of those in the poorest decile (the "poor"). However, it ignores information about incomes in the middle of the income distribution, and does not even use information about the distribution of income within the top and bottom deciles.

6.3 Inequality comparisons

Many of the tools used in the analysis of poverty can be similarly used for the analysis of inequality. In a way analogous to a poverty profile (see chapter 7), one could draw a profile of inequality, which among other things would look at the extent of inequality among certain groups of households. This provides information on the homogeneity of the various groups, an important element to take into account when designing policy interventions.

One may also analyze the nature of changes in inequality over time. One could focus on changes for different groups of the population to show whether inequality changes have been similar for all or have taken place, say, in a particular sector of the economy. In rural Tanzania, although average incomes increased substantially between 1983 and 1991, inequality increased (with the Gini coefficient increasing from 0.52 to 0.72), especially among the poor. This can be linked to important reforms that took place in

agricultural price policy, which intensified inequalities, with the poor and less-efficient farmers failing to participate in the growth experienced by wealthier, more efficient farmers (Ferreira, 1996).

It is often instructive to analyze other dimensions of inequality. For instance, in a country where public health provision is well developed and reaches all strata of the population, one could expect to see lower levels of inequality in health outcomes than in income levels, a proposition that could also be tested formally.

6.4 Decomposition of income inequality

The common inequality indicators mentioned above can be used to assess the major contributors to inequality, by different subgroups of the population and regions as well as by income source. For example, average income may vary from region to region, and this alone implies some inequality "between groups." Moreover, incomes vary inside each region, adding a "within group" component to total inequality. For policy purposes it is useful to be able to decompose these sources of inequality: if most inequality is due to disparities across regions, for instance, then the focus of policy may need to be on regional economic development, with special attention to helping the poorer regions.

More generally, in static decompositions, household and personal characteristics, such as education, gender, occupation, urban and rural, and regional location, are determinants of household income. If that is the case, then at least part of the value of any given inequality measure must reflect the fact that people have different educational levels, occupations, genders, and so on. This inequality is the "between-group" component.

But for any such partition of the population, whether by region, occupation, sector or any other attribute, some inequality will also exist among those people within the same subgroup; this is the "within-group" component. The Generalized Entropy class of indicators, including the Theil indexes, can be decomposed across these partitions in an additive way, but the Gini index cannot.

To decompose Theil's T index (i.e. GE(1)), let Y be the total income of the population, Y_j the income of a subgroup, N the total population, and N_j the population in the subgroup. Using T to represent GE(1)

(6.7)

$$T = \sum_{i=1}^{N} \frac{y_i}{N\bar{y}} \ln\left(\frac{y_i N}{\bar{y} N}\right) = \sum_{i=1}^{N} \frac{y_i}{Y} \ln\left(\frac{y_i N}{Y}\right)$$
$$= \sum_{j}\left(\frac{Y_j}{Y}\right) T_j + \sum\left(\frac{Y_j}{Y}\right) \ln\left(\frac{Y_j / Y}{N_j / N}\right)$$

This decomposes the inequality measure into two components. The first term represents the within-group inequality and the second term represents the between-group inequality. Similarly, GE(0) can also be decomposed. Using L to represent GE(0):

(6.8)

$$L = \sum_{i-1}^{N} \frac{1}{N} \ln\left(\frac{Y}{Y_i N}\right) = \sum_{j}\left(\frac{N_j}{N}\right) L_j + \sum_{j} \frac{N_j}{N} \ln\left(\frac{N_j / N}{Y_j / Y}\right).$$

Exercise: Decompose Theil's T measure of inequality into "within" and "between" components, using the income data provided in Table 6.2. [Hint: "Within" inequality should account for 69.1% of all inequality.]

For a typical decomposition of inequality in expenditure per capita, consider the following simple example, again from Dollar and Glewwe (1999, p.41), which refers to Vietnam in 1993. Using Theil's T, Table 6.4 shows that 22% of the total inequality is attributable to between-group inequality - i.e. to the difference in expenditure levels between urban and rural areas. The remaining 78% of all inequality is due to the inequality in expenditure per capita that occurs within each region.

Table 6.4: Decomposition of expenditure inequality by area, Vietnam, 1993

	Theil T	Between-group inequality	Memo: Population share (%)
All Vietnam	0.200		100
Urban only	0.196	0.044 (22% of total)	20
Rural only	0.136		80

Source: Dollar and Glewwe (1999), p.41.

Similar results were found for Zimbabwe in 1995-96. There a decomposition of Theil's T coefficient showed that the within-area (within rural areas and within urban areas) contribution to inequality was 72 percent, while the between–area (between urban and rural areas) component was 28 percent. In many

Latin American countries, the between-area component of inequality explains a much higher share of total inequality.

Of equal interest is which of the different income sources, or components of a measure of well-being, are primarily responsible for the observed level of inequality. For example, if total income can be divided into self-employment income, wages, transfers, and property income, one can examine the distribution of each income source. If one of the income sources were raised by one percent, what would happen to overall inequality?

Table 6.5 shows the results for the Gini coefficient for income sources in Peru (1997). As the table shows, self-employment income is the most equalizing income source. Thus a 1% increase in self-employment income (for everyone that receives such income) would lower the Gini by 4.9%, which represents a reduction in overall inequality. On the other hand, a rise in property income would be associated with an increase in inequality.

Generally, results such as these depend on two factors:

(1) the importance of the income source in total income (for larger income sources, a given percentage increase will have a larger effect on overall inequality), and

(2) the distribution of that income source (if it is more unequal than overall income, an increase in that source will lead to an increase in overall inequality).

Table 6.5 also shows the effect on the inequality of the distribution of *wealth* of changes in the value of different sources of wealth.

Table 6.5: Peru: Expected change in income inequality resulting from a one percent change in income source, 1997 (as percentage of Gini change)

Income source	Expected change	Wealth sources	Expected change
Self-employment income	-4.9	Housing	1.9
Wages	0.6	Durable goods	-1.5
Transfers	2.2	Urban property	1.3
Property income	2.1	Agricultural property	-1.6
		Enterprises	0

A final example, in the same spirit, comes from Egypt. There it was found that, in 1997, agricultural income represented the most important inequality-increasing source of income, while non-farm income has the greatest inequality-reducing potential. Table 6.6 sets out this decomposition and shows that while

agricultural income only represents 25% of total income in rural areas, it accounts for 40% of the inequality.

Table 6.6: Decomposition of income inequality in rural Egypt, 1997				
Income Source	Percentage of households receiving income from this source	Share in total income (%)	Concentration index for the income source	Percentage contribution to overall income inequality
Non-farm	61	42	0.63	30
Agricultural	67	25	1.16	40
Transfer	51	15	0.85	12
Livestock	70	9	0.94	6
Rental	32	8	0.92	12
All sources	100	100		100

References

Dollar, David and Paul Glewwe. 1999.

Ferreira. 1996.

Vietnam: General Statiistical Office. 2000. *Viet Nam Living Standards Survey 1997-1998*, Statistical Publishing House, Hanoi.

Chapter 7. Describing Poverty: Poverty Profiles

Summary

A poverty profile sets out the major facts on poverty and examines the pattern of poverty to see how it varies by

- Geography (by region, urban/rural, mountain/plain, etc.)
- Community characteristics (e.g. villages with and without a school, etc.)
- Household and individual characteristics (e.g. educational level).

A well-presented poverty profile can be immensely useful in assessing how economic change is likely to affect aggregate poverty, even though it typically uses rather basic techniques such as tables and graphs.

Some tables show the poverty rate for each group – e.g. by level of education of household head, or b region of the country. It is good practice to show the confidence interval of the poverty rates, which works especially well when the information is shown graphically. Alternatively one may show what fraction of the poor have access to facilities (e.g. running water, electricity) or live in a given region, and compare this to the non-poor. These are illustrated in the chapter by a number of graphs and tables based on data from Cambodia and Indonesia.

The World Bank's *Poverty Reduction* handbook has a long list of questions that a poverty profile should address. Provided that the data are available, it is very helpful to show how poverty has evolved over time. The change can often be linked to economic growth, and sometimes to specific government policies.

Most household surveys do not sample enough households to allow the analyst to break down the results at the sub-regional level. Yet poverty targeting – building roads, providing grants to poor villages, and the like – typically requires such detail. One solution is to use a poverty mapping: use the survey data to relate a household's poverty econometrically to a set of variables that are also available from the census; then apply the estimated regression equation to the census data in order to estimate whether a household is poor; this information can then be aggregated to give poverty rates for quite small areas.

A poverty profile is descriptive, but it serves as the basis for the analysis of poverty.

Learning Objective

After completing the module on poverty lines, you should be able to:

30. Explain what a poverty profile is and why it is useful.
31. Design tables and graphs that show the dimensions of poverty clearly and effectively.
32. Show why the use of additive poverty measures – such as the FGT class of measures (see Chapter 4) can facilitate poverty comparisions.
33. Explain why, in making poverty comparisons over time, one must correct for differences in sampling frame and method; adjust for price differences; and ensure comparability in the measures of income or expenditure.
34. Compute the relative risk of being poor for different household groups.
35. Summarize the steps required to undertake a poverty mapping, and explain why such a mapping is of practical value.

7.1 What is a country poverty profile?

A country poverty profile sets out the major facts on poverty (and typically, 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 household size). Hence, a poverty profile is simply a comprehensive poverty comparison, showing how poverty varies across sub-groups of society, such as region of residence or sector of employment. A well-presented poverty profile can be immensely informative and extremely useful in assessing how the sectoral or regional pattern of economic change is likely to affect aggregate poverty, even though it typically uses rather basic techniques such as tables and graphs.

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. Figure 7.1 shows that approximately 40% of the rural population, 10% of the Phnom Penh population and 25% of other urban residents live in households that are below the poverty line. Figure 7.1 also shows the 95% confidence interval that surround the estimates of the headcount index for each area. We interpret these confidence intervals to mean that we are 95% certain that the true poverty rate falls within these intervals. They reflect sampling error; other things being equal, the larger the sample, the narrower the confidence interval.

These standard error bands can be especially helpful when the sub-populations include only a small number of observations, because the bar charts may otherwise give a misleading sense of confidence in the precision of the poverty comparison that is illustrated. In the Cambodian case, the sampling errors are sufficiently small to have full confidence in the conclusion that headcount poverty rates are lower in Phnom Penh than in other urban areas, which in turn are lower than in rural areas. In terms of contribution to the total amount of poverty, 91% of people living below the poverty line live in rural areas, 7% live in other urban areas and 2% live in Phnom Penh, as the highlighted bars in Figure 7.1 show.

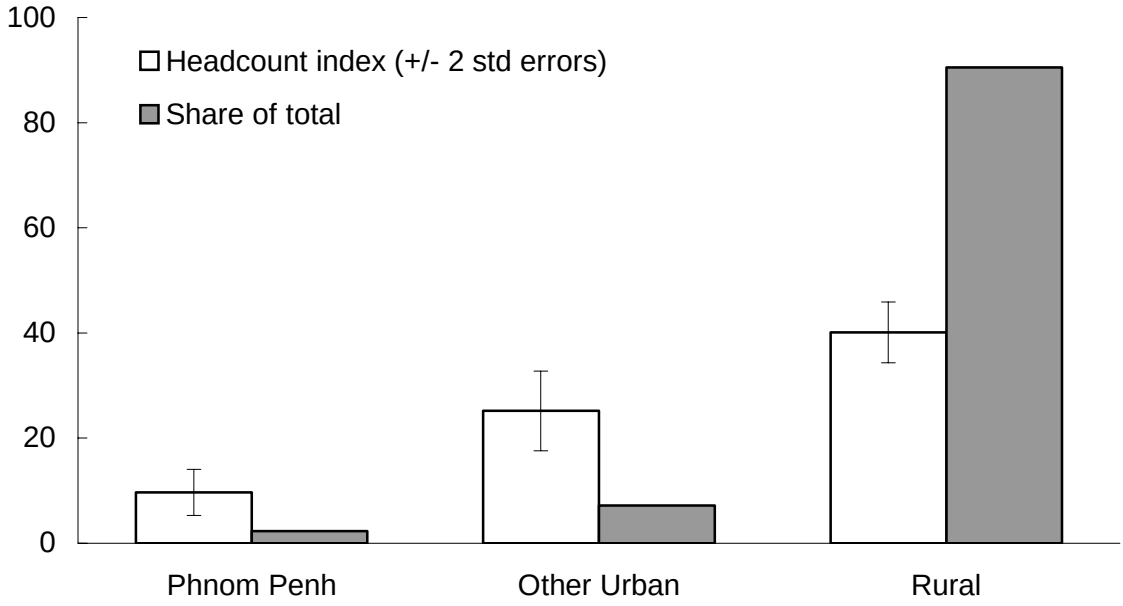

Figure 7.1: Cambodia: Headcount Poverty by Region

For our next example, Table 7.1 presents information on Ecuadorian households' access to services. The table shows, for instance, that 52% of the non-poor have waste collection, compared with just 24% for poor households. On average, the poor have lower access to services. An interesting finding, however, is that *within* urban areas, the poor have almost as much access to electricity as the non-poor; in this case essentially all the poor/non-poor differential occurs in rural areas. Note that we have rounded the figures to the nearest percentage point; this is to avoid giving an impression of spurious accuracy.

Table 7.1: Selected Characteristics of the Poor in Ecuador, 1994

Percentage with access to basic services:	Urban		Rural		Total	
	Poor	Non-poor	Poor	Non-poor	Poor	Non-poor
Sewerage connection	57	83	12	28	30	64
Electricity supply	98	100	62	76	76	91
Water from public net	61	79	18	23	35	59
Waste collection	60	77	1	6	24	52

Source: World Bank (1996)

By way of a further illustration, Table 7.2 shows poverty measures by household characteristics – gender and education levels of household head – for Malawi in 1997-98. Clearly, the higher the education level that household heads achieve, the less likely that the household is poor. This is a standard finding, but tables such as Table 7.2 help quantify the size of the effect.

Table 7.2: Poverty among household groups in Malawi, 1997-98

Household Characteristics	Headcount (P_0), %	Poverty gap (P_1)	Squared Poverty Gap (P_2)
Gender of Head			
Male	58	0.22	0.11
Female	66	0.28	0.15
Education levels of Head			
No education	71	0.31	0.17
Less than standard IV	63	0.25	0.13
Standard IV	58	0.22	0.11
Primary school	47	0.15	0.06
Secondary school	30	0.08	0.03
University	16	0.07	0.04

Source: National Economic Council (2000).

The World Bank's *Poverty Reduction* Handbook (1992) sets out some key questions that one may ask when preparing a poverty profile, as follows:

1. Does poverty vary widely between different areas in the country?
2. Are the most populated areas also the areas where most of the poor live?
3. How is income poverty correlated with gender, age, urban and rural, racial, or ethnic characteristics?
4. What are the main sources of income for the poor?
5. On what sectors do the poor depend for their livelihood?
6. What products or services—tradables and non-tradables—do the poor sell? A tradable good is one that is, or easily might be, imported or exported. The prices of such goods are influenced by changes in the world price and the exchange rate.
7. To what extent are the rural poor engaged in agriculture? In off-farm employment?
8. How large a factor is unemployment? Underemployment?
9. Which are the important goods in the consumption basket of the poor? How high is the share of tradables and non-tradables?
10. How is income poverty linked to malnutrition or educational outcomes?
11. What are fertility characteristics of the poor?
12. To what public services do the poor have access? What is the quality of these services?
13. How important are private costs of education and health for the poor?

14. Can the poor access formal or informal credit markets?
15. What assets—land, housing, and financial—do the poor own? Do property rights over such assets exist?
16. How secure is their access to, and tenure over, natural resources?
17. Is environmental degradation linked to poverty?
18. How variable are the incomes of the poor? What risks do they face?
19. Are certain population groups in society at a higher risk of being poor than others are? Households that are at a high risk of being poor, but are not necessarily poor now, are considered to be *vulnerable*.

A poverty profile that presents, in clear and readable form, answers to the above questions would be very helpful. But the extent to which a detailed poverty profile can be constructed depends on what data are available. While certain variables such as educational, health indicators, and access to essential services are the most basic components of a poverty profile, the relevance of many other variables depends on country circumstances. The general rule is that all variables that correlate with poverty and are relevant for policies under consideration should be included. By this rule, income generating activities, asset positions, access to social and infrastructure services, and the composition of consumption are all of interest. Cross-links to non-income measures of poverty may also be useful.

7.2 Additive poverty measures

The use of additive poverty measures can greatly facilitate poverty comparisons. Consider the general class of Foster, Greer and Thorbecke poverty measures described by equation 4.6 of chapter 4. Suppose the population can be divided into m mutually exclusive sub-groups. The poverty profile is simply the list of poverty measures P_j for j=1, … ,m. Aggregate poverty can then be written as the average of the sub-group poverty measures, weighted by their population:

(7.1)
$$P = \frac{1}{N}\sum_{j=1}^{m} P_j N_j$$

where

(7.2)
$$P_j = \frac{1}{N_j}\sum_{i=1}^{N_j} p(z_j, y_i^j)$$

is the poverty measure for the sub-group j with population N_j. Here y_i^j is the welfare indicator of individual i who belongs to the sub-group j, where i=1, … ,N_j. The total population N is equal to $\sum_{j=1}^{m} N_j$

Analogously, one can also define "clusters" of sub-groups; as one disaggregates further and further, the poverty profile at each step adds up to that of the previous step, using population weights.

In addition to the computational convenience in forming poverty profiles, additive poverty measures guarantee "subgroup consistency" in the sense that when poverty increases in any sub-group of the population, aggregate poverty will also increase, other things being equal. It can be shown that, given certain assumptions of technical nature, subgroup consistency implies, and is implied by, the class of measures defined by equation 4.6.

It is sometimes objected that additivity attaches no weight to the differences between sub-groups in the extent of poverty. Consider two equal-sized groups – "rural" and "urban" sectors – with initial headcount poverty indices of 70% and 20% respectively. Aggregate poverty is 45% according to any (population weighted) additive measure. Now you are asked to choose between two policies X and Y. Under policy X, the poverty profile changes to 70% and 10%, while under policy Y the profile becomes 60% and 20%. By any additive poverty measure one should be indifferent between X and Y; both yield an aggregate poverty index of 40%. Yet, in contrast to X, the gains under policy Y have gone to the poorer rural sector. Should we prefer policy Y? The answer is only "yes" if one is concerned about inequalities between groups independently of absolute living standards, and it is by no means clear why we should be.

7.3 Profile Presentation

There are two main ways of presenting a poverty profile. The first ("type A") gives the incidence of poverty or other poverty measure(s) for each sub-group defined in terms of some characteristic, such as place of residence, as is done in Table 7.2. The second ("type B") gives the incidence of characteristics amongst sub-groups defined in terms of their poverty status, such as "poor" and "non-poor", as is done in the highlighted columns in Figure 7.1.

Which is more useful will depend on the purpose of the poverty profile. Suppose that one is using the poverty profile to select a target region for a poverty alleviation scheme. The scheme will allocate a small sum of money to all residents in the chosen target region ("indicator targeting"). It is an

imperfect form of targeting because (as is invariably the case) the policymaker does not know who has which standard of living, even when a distribution of living standards can be constructed from a household sample survey. The policymaker therefore has to make do with an imperfect indicator of living standards, in this case region of residence. In this case a type A profile will be more useful, because it identifies the areas (or groups) where poverty rates are highest, and hence where universal cash transfers are least likely to be handed out to the non-poor.

7.4 Poverty Comparisons over time

If several consecutive rounds of household surveys are available, changes in income poverty over time can be assessed. This requires poverty measures that are comparable and that reflect differences over time in the cost of living across regions. The commonest method for preparing comparisons over time consists of converting nominal income or consumption data from different surveys and regions into real income and consumption by deflating the indicators in space and time. A constant poverty line can then be applied to these real values to infer poverty measures. Ideally, to obtain robust poverty comparisons over time, one would want to use surveys with similar sampling frame and methods, with corrections for prices differences, and with similar definitions of consumption or income. In practice, however, differences exist in some of these dimensions. This does not imply that no comparison can be made. It simply means that the analyst will need to:

- correct for major differences in the sampling frame and sampling method for the different surveys or the different rounds of a panel survey;
- use regional and temporal price indices to ensure a similar definition of the poverty line over time and across regions;
- adjust the definition of consumption or income aggregates over time to ensure that a similar definition is used. However, one should be aware that changes in definitions, and in particular in the degree to which home production is included in the definition, can lead to significant distortions of poverty measurement.

When several rounds of survey data are available, the analyst can investigate changes in the regional distribution of poverty or in the major characteristics of the poor, such as ethnicity, gender, age, urban and rural location, employment, access to social programs and basic services, etc. Although the various population groups identified in the first period of time will most naturally form the basis of the analysis over time, it is also important to investigate whether or not 'new' groups of poor people have

appeared. This is particularly relevant for countries that undergo rapid changes linked to such factors as economic reforms, conflicts, natural disasters, and epidemics such as HIV/AIDS.

Table 7.3 compares the baseline poverty profile for Cambodia derived from the 1993/94 data with that of the CSES 1997. Note that the value of the poverty line (consisting of the food poverty line plus a non-food allowance equal to the level of non-food consumption of persons whose per capita consumption just equals the food poverty line following Ravallion's model) increased by 15% in Phnom Penh, 11% in other urban areas, and 8% in rural areas.

Table 7.3: Poverty measures for Cambodia, 1993-94 and June 1997

	Head count index (P_0), %		Poverty gap index (P_1), %		Poverty severity index (P_2), ×100		Memo: Poverty line, riels/day	
	1993/94	1997	1993/94	1997	1993/94	1997	1993/94	1997
Food Poverty Line								
Phnom Penh	6.2	3.4	1.3	0.5	0.4	0.1	1,578	1,819
Other Urban	19.6	15.4	4.4	3.3	1.4	1.1	1,264	1,407
Rural	21.9	20.0	4.0	3.9	1.1	1.2	1,117	1,210
Total	20.0	17.9	3.7	3.5	1.1	1.1		
Poverty Line								
Phnom Penh	11.4	11.1	3.1	2.2	1.2	0.6		
Other Urban	36.6	29.9	9.6	7.5	3.6	2.7		
Rural	43.1	40.1	10.0	9.7	3.3	3.4		
Total	39.0	36.1	9.2	8.7	3.1	3.1		

Source:

The estimates in table 7.3 indicate that the incidence of poverty declined modestly in Cambodia as a whole (from 39% to 36%) during the period 1993/94 to June 1997. On a regional basis, poverty declined significantly in other urban areas (from 37% to 30%), modestly in rural areas (from 43% to 40%) and not at all in Phnom Penh (where it remained at 11%). During the same period, the estimates indicate that two other measures of poverty (i.e. the poverty gap and poverty severity index) declined significantly, both in Phnom Penh and in other urban areas but not in rural areas.

Poverty measures are sometimes translated into the *relative risks* of being poor for different household groups. These risks indicate whether the members of a given group are poor in relation to the corresponding probability for all other households of society. This concept may be applied to examine whether, over time, the relative poverty risk of specific population groups decreases or increases. Table 7.4 compares the relative poverty risk of various groups in Peru in 1994 and 1997. It shows, for instance,

that households with seven persons or more were 71% more likely to be poor in 1994 than other households in society; and that this relative risk was 106% in 1997 (i.e. they were more than twice as likely to be poor as other households in Peru). Or again, between 1994 and 1997 the relative risk of being poor for households where the spouse of the head was working diminished (from –11% to –21%).

Table 7.4: Poverty risks for selected groups of households, Peru, %		
Household characteristic	1994	1997
Households using house for business purposes	-28	-29
Rural households with at least one member in off-farm employment	-24	-23
Households where head of spouse was working	-11	-21
Households without water or sanitation	54	50
Households without electricity	63	69
Households where head had less than secondary education	73	72
Households of 7 persons or more	71	106
Source: ***xxx***		

7.5 Excerpts from poverty profiles for Indonesia and Cambodia

In this section we present some excerpts from poverty profiles for Indonesia and Cambodia. These give a flavor of the types of tables and figures that are typically constructed for poverty profiles, and that are well worth imitating!

7.5.1 Indonesia

Table 7.5 gives an example of a poverty profile in which the sampled households in Indonesia's 1987 SUSENAS have been classified into eleven groups according to their principal income source. Results are given for the three main poverty measures discussed above. The following points are worth noting:

- In the absence of adequate information on urban versus rural prices, Huppi and Ravallion (1991) assumed an urban-rural cost-of-living differential of 10%. Although this appears to be a reasonable assumption, their results are sensitive to this assumption.
- The poverty measures are based on the estimated population distributions of persons ranked by household consumption per person, where each person in a given household is assumed to have the same consumption. Household specific sampling rates have been used in estimating the distributions.
- In forming the poverty profile, households have been grouped by their stated "principal income source." Many households will have more than one income source. In principle one could form sub-groups according to the various interactions of primary and secondary income sources, but this would

rapidly generate an unwieldy poverty profile; as a general rule, it is important to keep poverty profiles relatively straightforward and uncluttered.

- The three measures are in close agreement on the ranking of sectors in terms of poverty. For example, the two farming sub-groups are the poorest by all three measures.

Table 7.5: Sectoral poverty profile for Indonesia, 1987

Principal sector of employment	Population share (1987)	Head-count index (P_0), %	Poverty gap index (P_1), %	Poverty severity index (P_2), ×100
Farming				
Self-employed	41.1	31.1	6.42	1.97
Laborer	8.6	38.1	7.62	2.21
Industry				
Urban	3.0	8.1	1.26	0.32
Rural	3.4	19.4	3.00	0.76
Construction	4.3	17.4	2.92	0.80
Trade				
Urban	6.3	5.0	0.71	0.17
Rural	7.6	14.7	2.42	0.61
Transport	4.1	10.7	1.53	0.34
Services				
Urban	7.6	4.2	0.61	0.14
Rural	7.3	11.6	1.84	0.49
Other	6.7	17.1	3.55	1.03
Total	**100.0**	**21.7**	**4.22**	**1.24**

Source: Huppi and Ravallion (1991).

Changes in the poverty profile may arise from the contributions of different sub-groups in the poverty profile to changes over time in aggregate poverty. Table 7.6 provides information on the relative contribution of various sectors to aggregate poverty alleviation in Indonesia between 1984 and 1987. These are the "intra-sectoral effects", expressed as a percentage of the reduction in aggregate poverty for each poverty measure. For instance, 11% of the reduction in poverty (as measured by P_0) between 1984 and 1987 was due to the fall in poverty among farm laborers. The table also gives the aggregate contribution of shifts in population and the interaction effects between sectoral gains and population shifts.

The drop in poverty among self-employed farmers had the largest influence on aggregate poverty reduction, and most particularly on the reduction in the severity of poverty as measured by P_2. About 50% of the reduction in the national head-count index was due to gains in this sector, while it accounted for 57% of the gain in P_2. Note that the rural farm sector's impressive participation in the reduction of aggregate poverty is due to both significant declines in their poverty measures, and the large share of national poverty accounted for by this sector.

Furthermore, 13% of the decline in the national head-count index was due to population shifts between various sectors of employment – essentially as people moved oute of high-poverty into low-poverty sectors. The sectors that gained in population share were almost all urban (Huppi and Ravallion 1991), and had initially lower poverty measures. The fact that population was moving out of the rural sector, where poverty was falling faster, accounts for the negative interaction effects in Table 7.6.

Table 7.6: Sectoral decomposition of the change in poverty in Indonesia, 1984-87

		Contribution of sectoral change in:		
Principal sector of employment	Population share (1984)	Head-count index (P_0), %)	Poverty gap index (P_1), %	Poverty severity index (P_2), ×100
Farming				
Self-employed	45.0	49.8	54.6	57.4
Laborer	9.0	11.2	14.8	16.5
Industry				
Urban	2.6	0.8	0.4	0.3
Rural	3.3	2.8	3.1	2.7
Construction	4.1	3.2	2.6	2.2
Trade				
Urban	5.4	2.2	1.6	1.4
Rural	6.6	7.2	5.6	4.7
Transport	3.8	3.6	2.7	2.2
Services				
Urban	6.5	1.0	1.0	0.9
Rural	5.8	2.9	2.4	2.0
Total sector effects (incl. omitted sectors)		89.3	93.8	95.1
Contribution of population shifts		13.2	10.4	9.4
Interaction effects		-2.6	-4.3	-4.5
Total	**100.0**	**100.0**	**100.0**	**100.0**

Source: Adapted from Huppi and Ravallion (1991).
Note: minor sectors omitted.

7.5.2 Cambodia

A basic breakdown of Cambodian poverty rates, by region, is given above in Table 7.3. The numbers show that at least 85% of the poor are concentrated in rural areas.

Some more recent figures are shown in table 7.7, using data from the CSES 1999. Data in 1999 were collected in two rounds, and table 7.7 contains estimates for each round (and the pooled sample) of the three main poverty statistics, and also reports the results from the previous surveys for comparison.

Table 7.7: Comparisons of poverty estimates from Cambodian surveys			
	Head-count index (P_0), %)	Poverty gap index (P_1), %	Poverty severity index (P_2), ×100
SESC 1993/94	39.0	9.2	3.1
1997 CSES (as adjusted by Knowles)	36.1	8.7	3.1
1997 CSES (unadjusted)	47.8 (1.5)	13.7 (0.7)	5.3 (0.3)
CSES 1999 (Round 1)	64.4 (2.3)	23.9 (1.3)	11.3 (0.8)
CSES 1999 (Round 2)	35.9 (2.4)	6.5 (0.7)	2.0 (0.4)
CSES 1999 (both Rounds combined)	51.1 (1.8)	15.4 (0.9)	6.7 (0.5)
Note: No sampling errors (reported in parentheses for the other years) are reported by the two previous poverty profiles but the relative errors for SESC 1993/94 and the adjusted 1997 CSES would likely be higher than the relative error in 1999 because the sampling scheme used previously was not as efficient (fewer clusters and broader stratification). The poverty line used for the unadjusted 1997 CSES results takes values of 1923 riels per person per day in Phnom Penh, 1398 in other urban and 1195 in rural.			

Source:

An interesting feature of these results is that there is a substantial discrepancy in the poverty estimates from the two survey rounds in 1999. The headcount index is almost 30 percentage points higher for Round 1 than for Round 2, while the poverty gap and poverty severity indexes are between four and six times higher. These troubling discrepancies are also large relative to the variation across previous survey estimates of poverty in Cambodia, and would need to be investigated and fully discussed in a serious poverty profile. If the discrepancies between the two survey rounds are ignored, and the data are pooled, the resulting poverty estimates are fairly similar to the unadjusted 1997 estimates, showing a slight *increase* in all three poverty measures (Table 7.7).

The pattern of poverty with respect to the age group of the household head is reported in Table 7.8. It is apparent that poverty rates rise with age, reaching a maximum for the 36-40 year old group of household heads, and then declining. A similar pattern was reported in the 1997 poverty profile. Once again, the definition of headship and its economic interpretation may confound the results so more detailed examination would be needed before any interventions might be designed on the basis of these age patterns. For example, the household head need not be the major economic contributor to the household; respondents may simply have nominated the oldest or most senior member. Thus, the relatively low poverty rate for people living in households whose head is aged 61 years and above may reflect the wealth accumulation that this elderly head has achieved or it could be that there is a younger generation within the household whose economic success is sufficient to allow them to support their elders within the same household. As a general rule, it is wise not to put too much emphasis on breakdowns by household head, given the problems involved in its definition. Reflecting this, the US

Census no longer even asks who is the head of the household; it has also become less socially acceptable to identify a "head" of household in the United States.

Table 7.8: Distribution of poverty by age and gender of household head

	Head-count index (P_0), %		Poverty gap index (P_1), %		Poverty severity index (P_2), ×100		Share of total pop (%)
	Index (%)	Contribution to total (%)	Index (%)	Contribution to total (%)	Index (%)	Contribution to total (%)	
Poverty Line	35.9	100.0	6.5	100.0	2.0	100.0	100.0
Age of head:							
18-30 years	36.7	10.7	5.6	9.1	1.4	7.5	10.5
31-35 years	35.4	10.9	5.4	9.2	1.6	8.8	11.1
36-40 years	43.6	21.2	8.0	21.6	2.7	23.3	17.5
41-45 years	40.3	15.7	7.3	15.8	2.2	15.3	14.0
46-50 years	36.5	14.4	7.7	16.9	2.4	16.9	14.2
51-60 years	28.3	15.8	5.3	16.3	1.7	16.8	20.0
61 and Above	32.0	11.3	5.6	11.1	1.8	11.3	12.7
Male	36.4	84.4	6.6	84.2	2.1	85.1	83.3
Female	33.6	15.7	6.1	15.8	1.8	14.9	16.7

Source: Gibson (1999).

Note that the poverty level is *lower* among female-headed households. This is not unusual in Southeast Asia. Often a finer breakdown is more helpful – for instance, households headed by widows, by married women with an absent husband (who may send remittances home), and so on.

There are two reasons why widow-headed households, and households where there has been a dissolution (i.e., separation or divorce), could be at greater risk of poverty. The loss of an economically active household member, as would occur with the death of a husband in war for example, is likely to cause a large income shock that could push a household into poverty. The second factor, and the one that links marital status with household size, is that widow-headed households tend to be smaller than average, which will constrain the effective living standards of their residents if there are economies of scale in household consumption.

In Cambodia, the headcount poverty rate in 1999 increased smoothly with household size to a maximum rate for households with eight members (Figure 7.2). In the Round 1 data, the highest headcount poverty rate was for households with nine members. A relationship like that shown in Figure 7.2 might normally be doubted because it does not control for economies of scale in household consumption: large households may have lower expenditures (per capita) not because their members are poor but because they do not need to spend as much per person to reach the same standard of living. However, there is some evidence that such economies of size are relatively unimportant for Cambodian households, in which case the pattern shown by Figure 7.2 may be a useful basis for identifying the poor.

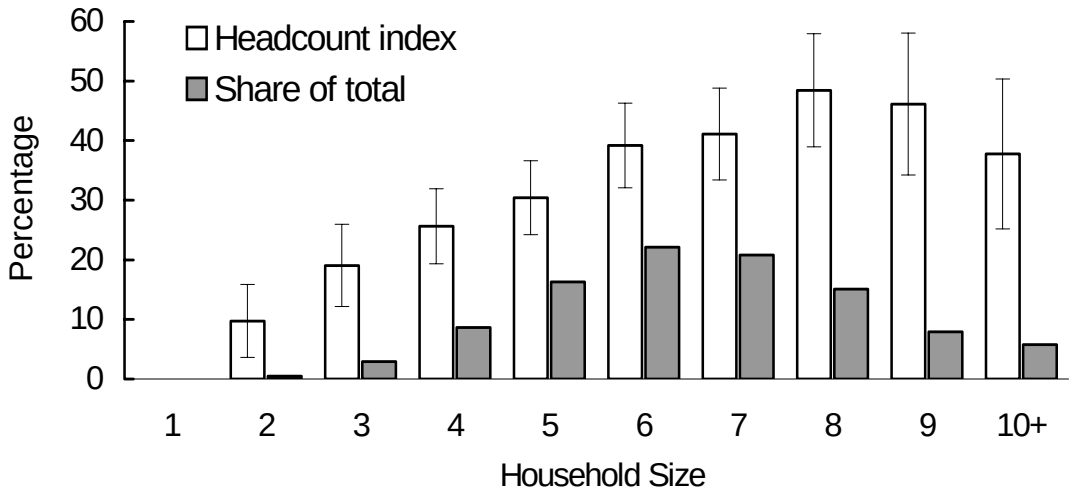

Figure 7.2: Poverty by household size

Previous poverty profiles showed that poverty rates were relatively high among those whose household head either has no schooling or has only primary schooling. Poverty rates then fall with the attainment of lower secondary education, fall farther with upper secondary and are almost zero if the household head is a university graduate. But those whose household head had a technical/vocational or other form of education had a higher poverty rate than those with primary schooling (at least in the 1997 poverty profile), for reasons that are not entirely clear. This is a good example of a case where the poverty profile raises questions that require further examination.

According to the survey estimates, there was no difference, in 1999, in poverty rates between those whose household head has no schooling and those whose head has some primary education (figure 7.3). Although the survey estimate of the headcount poverty rate is slightly higher for the primary schooled group, the estimates for both groups are surrounded by wide and overlapping confidence intervals. One possible explanation for this somewhat surprising result is that primary education is of very low quality, so it adds little to one's earning ability. The finding is in line with evidence from a number of other countries that suggests that a secondary education is required to truly pull someone out of poverty,.

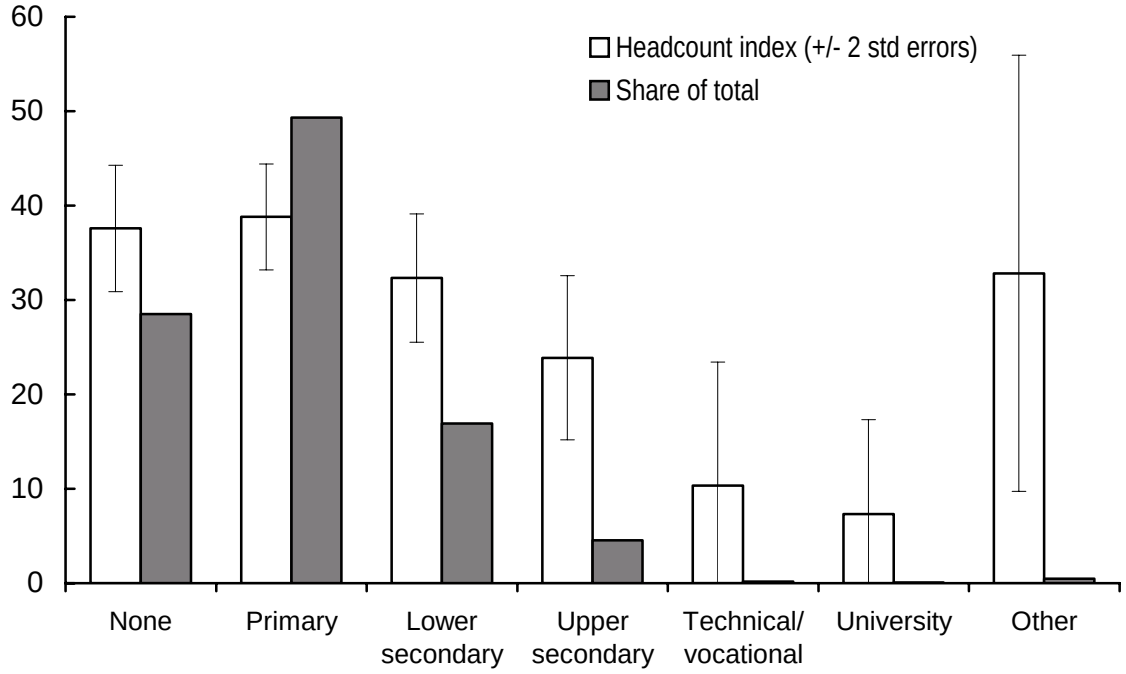

Figure 7.3: Poverty by education level of household head

7.6 Poverty mapping

Poverty analysis is often based on national level indicators that are compared over time or across countries. The broad trends that can be identified using aggregate information are useful for evaluating and monitoring the overall performance of a country. For many policy and research applications, however, the information that can be extracted from aggregate indicators may not be sufficient, since they hide significant local variation in living conditions within countries. For example, poverty within a region can vary across districts. This makes small-area estimates of poverty very appealing. However, often we are unable directly to compute poverty estimates for small areas like districts. Instead, we usually have poverty estimates for regions or entire countries only.

The main reason that poverty measures are computed for large areas and not usually available for small areas is data availability. There are two main types of welfare-related information sources available to policy-makers. *Household surveys* often include a detailed income and/or consumption expenditure module (such as the Cambodian CSES 1999). However, due to the relatively small sample size, the collected information is usually only representative for broad regions of the country. For example, with

the CSES 1999 we can compute poverty estimates for Phnom Penh, other urban areas, and rural areas, but not for every district. *Census data* (and sometimes large household sample surveys) are available for all households (or very large samples of households) and can provide reliable estimates at highly disaggregated levels such as small municipalities, towns, and villages. But censuses do not contain the income or consumption information necessary to yield reliable indicators of the level and distribution of welfare such as poverty rates or inequality measures.

However, it is possible to merge information from these two types of data sources (detailed household surveys like the CSES 1999, and census data) so that "poverty maps" can be constructed; Box 7.1 provides some further details. These detailed poverty maps capture the heterogeneity of poverty within a country. That is, areas that are better off and those that are worse off will be more clearly defined. Sometimes regions that have less poverty overall may have substantial pockets of poverty that are lost in the aggregate poverty statistics.

Poverty maps can improve the targeting of interventions. In designing poverty alleviation projects and allocating subsidies, resources will be used more effectively if the most-needy groups can be better targeted. This reduces the leakage of transfer payments to non-poor persons, and it reduces the risk that poor persons will be missed by a program. Poverty maps can also help governments to articulate their policy objectives. Basing allocation decisions on observed geographic poverty data rather than subjective rankings of regions increases the transparency of government decision-making. Such data can thus help limit the influence of special interests in allocation decisions. There is a role for well-defined poverty maps in lending credibility to government and donor decision-making.

To create a poverty map,

1. First, use the household survey data to estimate a model of per capita consumption expenditure (or any other household or individual-level indicator of well-being) as a function of variables that are common to both the household survey and the census. Such variables typically include household size, the educational level of the household head, the proportion of the household consisting of prime-age adults, and sometimes information about the quality of the housing.
2. Second, use the resulting estimated equation to predict per capita expenditures for each household in the census.
3. Third, use the estimated household-level measures of expenditure per capita to determine which households are poor.

4. Fourth, aggregate this information to the level desired – for instance, village or commune or parish or municipality.

For a recent application to Vietnam, see Baulch and Minot (2001); discussions of the methodology may be found in Hentschel et al. (2000), Elbers, Lanjouw and Lanjouw (2000), and Alderman et al. (2000).

Box 7.1: Poverty Mapping In Ecuador

In this box we discuss a poverty mapping of Ecuador – how it was done, and what results emerged. The discussion rests heavily on Hentschel et al. (2000). Ecuador has about 400 cantons and over 1,000 parishes (*parroquias*). The purpose of the exercise was to get poverty rates for each parish in the country. In principle this would allow a relatively finely-tuned targeting of the poorer parts of the country.

The 1994 *Encuesta sobre las Condiciones de Vida* (a Living Standards Measurement Survey – see http://www.worldbank.org/html/prdph/lsms/lsmshome.html) obtained 4,391 usable responses from households, which was clearly inadequate for measuring poverty at the level of each parish or even canton. However the 1990 census counted about 2 million households, and collected information on a range of demographic variables such as household size, age, education, occupation, housing quality, language and location

The research team used the data from the *Encuesta* to estimate regressions of the form

$$\ln y_i = \mathrm{X}_i' b + e_i$$

for each region of the country. The dependent variable was income per capita, and the independent variables included such measures as the age, gender and education of the head of the household, which were available both from the *Encuesta* and also from the census. With R^2 values of about 0.5, the fits were adequate. Then data from the census for each household were then inserted into the equation in order to predict the level of income for each household, and poverty rates were computed for each parish.

The first finding was that the poverty rates for each of the broad regions are robust. The authors re-estimated the income equation using part of the *Encuesta* sample, predicted income for all households that were not included in the estimation (the validation sample), and then compared the predicted income with actual income. These out-of-sample predictions proved to be quite close to the actual values. However, the poverty rates by parish were not robust, as the standard errors of the estimates were relatively high.

In another test, Hentschel at al. simulated the effect of providing subsidies to the poorest parishes, and then asked what proportion of these benefits went to each income quintile. The results are reproduced here, and show that 78% of the subsidies would have gone to households in the lowest two quintiles – a respectably high level of successful targeting.

Distribution of beneficiary households with geographic targeting at the parish level

	Income quintile				
	Lowest	Low-mid	Middle	Mid-upr	Upper
% of beneficiary households	51.0	27.0	13.1	8.0	0.9

Source: Hentschel et al.

Their conclusion? "The most useful practical application of this methodology is probably in making comparisons with regional patterns of other indicators of well-being, opportunity, and access" (Hentschel et al., p.162). Thus, for instance, one could map health indicators against estimated income at the regional level in order to investigate a possible close and useful link.

References

Alderman et al. (2000).

Baulch and Minot (2001);

Elbers, Lanjouw and Lanjouw (2000),

Gibson? 1999. Cambodia data.

Hentschel, Jesko, Jean Olson Lanjouw, Peter Lanjouw and Javier Poggi. 2000. "Combining Census and Survey Data to Trade the Spatial Dimensions of Poverty: A Case Study of Ecuador," *The World Bank Economic Review*, 14(1): 147-166.

Huppi and Martin Ravallion. 1991.

Malawi, National Economic Council.

Peru data

World Bank. 1992. *Poverty Reduction Handbook.* Washington DC.

World Bank. 1996. [Ecuador]

Nicholas Prescott and Menno Pradhan. 1997. *A Poverty Profile for Cambodia*, World Bank, Washington DC.

Nguyen Binh Duong and Dinh Hien Trinh. 1999. "Identification of Poverty in Vietnam," in Dominique and Jonathan Haughton et al. (eds.), *Health and Wealth in Vietnam: An Analysis of Household Living Standards*, Institute of Southeast Asian Studies Press, Singapore.

Vu Trieu Minh. 1999. "Household Expenditure," in Haughton et al., *Health and Wealth in Vietnam.*

Chapter 8. Understanding the Determinants of Poverty

Summary

A poverty profile describes the pattern of poverty, but is not principally concerned with explaining its causes. Yet a satisfactory explanation of why some people are poor is essential if we are to be able to tackle the roots of poverty.

Among the key causes, or at least correlates, of poverty are:

Regional-level characteristics: these include vulnerability to flooding or typhoons; remoteness; quality of governance; property rights and their enforcement.

Community level characteristics: these include the availability of infrastructure (roads, water, electricity) and services (health, education), proximity to markets, and social relationships.

Household and individual characteristics: Among the most important are:

Demographic: household side, age structure, dependency ratio, gender ofhead.

Economic: employment status, hours worked, property owned.

Social: health and nutritional status, education, shelter.

Regression analysis is commonly undertaken to identify the effects of each of these characteristics on income (or expenditure) per capita. Care is needed to choose the independent variables carefully, to be sure that they are indeed exogenous. A number of more exotic techniques are now available for this purpose, including CART (classification and regression tree) models and multiple-adaptive regression splines (MARS models).

Regression techniques are good at identifying the immediate ("proximate") causes of poverty, but are less successful at finding the "deep" causes; they can show that a lack of education causes poverty, but cannot so easily explain why some people lack education.

Learning Objectives

After completing the module on poverty lines, you should be able to:

36. Identify the main immediate ("proximate") causes of poverty.
37. Classify the main causes of poverty into characteristics related to the country or region, the community, and the household and individual.
38. Explain how regression techniques may be used to identify the proximate causes of poverty and their relative importance.
39. Explain why researchers generally prefer to use regressions to "explain" income (or expenditure) per capita rather than whether an individual is poor or not.
40. Evaluate the assertion that the weakest part of poverty analysis is the understanding of its fundamental causes, and that this represents a "missing middle" that makes it difficult to define a successful anti-poverty strategy.

8.1 What causes poverty?

A poverty profile describes the pattern of poverty, but is not principally concerned with explaining its causes. Yet a satisfactory explanation of why some people are poor is essential if we are to be able to tackle the roots of poverty. This chapter addresses the question of what causes poverty.

Poverty may be due to national, sector-specific, community, household or individual characteristics. This chapter summarizes some of the characteristics of the poor by region, community, household and individual characteristics and then discusses how regression techniques can be used to determine the factors "causing" poverty.

Two cautions are in order. First, it can be difficult to separate causation from correlation. For instance, we know that poor people tend to have low levels of education; but are they poor because they have little education, or do they have little education because they are poor? A statistical association alone is not enough to establish causality, and additional information is likely to be required.

Second, most of the "causes" of poverty that we identify in this chapter are immediate (or "proximate") causes, but not necessarily "deep" causes. For instance, suppose that we can demonstrate that low levels of education do indeed increase the risk of poverty. This is interesting, but now begs the question of why some people have low levels of education in the first place: Were the school fees too high? Was there no school nearby? Was the quality of the education abysmal? Were their parents unsupportive, or even hostile to education? Was there a concern that an educated woman could not find a husband?

The weakest part of poverty analysis – what Howard White calls the "missing middle" – is developing a clear understanding of the fundamental causes of poverty. Such an understanding is needed if one is to develop an effective strategy to combat poverty. Since there is no reason to believe that the root causes of poverty are the same everywhere, country-specific analysis is essential.

8.1.1 Regional level characteristics

At the regional level, there are numerous characteristics that might be associated with poverty. The relationship of these characteristics with poverty is country-specific. In general, however, poverty is high in areas characterized by geographical isolation, a low resource base, low rainfall, and other inhospitable climatic conditions. For example, many argue that economic development in Bangladesh is severely retarded due to its susceptibility to annual floods; and Nghe An province in north-central

Vietnam is poor in part because it is regularly hit by typhoons, which destroy a significant part of the accumulated stock of capital. In many parts of the world the remoteness of rural areas – which lower the price farmers get for their goods and raise the price they pay for purchases, due to high transport costs – is responsible for generating food insecurity among the poor. Inadequate public services, weak communications and infrastructure, as well as underdeveloped markets are dominant features of life in rural Cambodia, as in many other parts of the world, and clearly contribute to poverty.

Other important regional and national characteristics that affect poverty include good governance, a sound environmental policy, economic, political and market stability, mass participation, global and regional security, intellectual expression and a fair, functional, and effective judiciary. Regional-level market reforms can boost growth and help poor people, but it is important to note that they can also be a source of dislocation. The effects of market reforms are complex, deeply linked to institutions and to political and social structures. The experience of transition, especially in countries of the former Soviet Union, vividly illustrates that market reforms in the absence of effective domestic institutions can fail to deliver growth and poverty reduction. There is also a case for bringing vulnerability and its management to center stage. Participatory poverty work underlines the importance of vulnerability to economic, health, and personal shocks.

Inequality is also back on the agenda. New work shows the importance of gender, ethnic, and racial inequality as a dimension – and a cause – of poverty. Social, economic, and ethnic divisions in regions are often sources of weak or failed development. In the extreme, vicious cycles of social division and failed development erupt into internal conflict (within or across regions), as in Bosnia and Sierra Leone, with devastating consequences for people.

8.1.2 Community level characteristics

As with regional characteristics, there are a variety of community-level characteristics that may be associated with poverty for households in that community. At the community level, infrastructure is a major determinant of poverty. Indicators of infrastructure development that have often been used in econometric exercises include proximity to paved roads, whether or not the community has electricity, proximity to large markets, availability of schools and medical clinics in the area, and distance to local administrative centers. Other indicators of community level characteristics include average human resource development, access to employment, social mobility and representation, and land distribution.

Recent research has also stressed the importance of social networks and institutions, and "social capital" (which includes, for instance, the level of mutual trust in the community) xxx: Robert Putnam. In addition to removing social barriers, effective efforts to reduce poverty require complementary initiatives to build up and extend the social institutions of the poor. Social institutions refer to the kinship systems, local organizations, and networks of the poor and can be thought of as different dimensions of social capital. Research on the roles of different types of social networks in poor communities confirms their importance. An analysis of poor villages in North India, for example, shows that social groups play an important role in protecting the basic needs of poor people and in mediating risk. Studies of agricultural traders in Madagascar show that social relationships are more important to traders than input prices. Close relationships with other traders are used to lower the transactions costs of exchange, while ties to creditors and others who can help out during times of financial hardship are vital sources of security and insurance.

How does social capital affect development? The narrowest view holds social capital to be the social skills of an individual – one's propensity for cooperative behavior, conflict resolution, tolerance and the like. A more expansive meso view associates social capital with families and local community associations and the underlying norms (trust, reciprocity) that facilitate coordination and cooperation for mutual benefit. A macro view of social capital focuses on the social and political environment that shapes social structures and enables norms to develop. This environment includes formalized institutional relationships and structures, such as government, the political regime, the rule of law, the court system, and civil and political liberties. Institutions have an important effect on the rate and pattern of economic development.

An integrating view of social capital recognizes that micro, meso, and macro institutions coexist and have the potential to complement one another. Macro institutions can provide an enabling environment in which micro institutions develop and flourish. In turn, local associations help sustain regional and national institutions by giving them a measure of stability and legitimacy – and by holding them accountable for their actions. Social capital is clearly a complicated characteristic and often researchers find it difficult to identify appropriate variables that measure social capital quantitatively.

8.2 Household and individual level characteristics

Some of the important characteristics in this category would include the age structure of household members, education, gender of the household head, and the extent of participation in the labor force. In

recent times, other components that fall under this category have included domestic violence prevention, and gender-based, anti-discrimination policies. The following discussion organizes these characteristics into groups and discusses them in greater detail. These groups are demographic, economic and social characteristics.

8.2.1 Demographic characteristics

Indicators of household size and structure are important in that they show a possible correlation between the level of poverty and household composition. Household composition, in terms of the size of the household and characteristics of its members (such as age), is often quite different for poor and non-poor households. The Cambodian CSES of 1993/94 shows that the poor tend to live in larger households, with an average family size of 6.6 persons in the poorest quintile compared to 4.9 in the richest quintile; similar patterns are found in most countries. The poor also tend to live in younger households – with the bottom quintile having twice as many children under 15 per family as the top quintile – and slightly fewer elderly people over age 60. Better-off households also tend to have heads that are somewhat older.

The dependency ratio is calculated as the ratio of the number of family members not in the labor force (whether young or old) to those in the labor force in the household. This ratio allows one to measure the burden weighing on members of the labor force within the household. One might expect that a high dependency ratio will be associated with greater poverty.

It is widely believed that the ***gender of the household head*** significantly influences household poverty, and more specifically that households headed by women are poorer than those headed by men. This is of particular importance to Cambodia. Due to male casualties in past wars, women are often the heads of households. Women play an important role in the labor force, both in the financial management of the household and in the labor market, but appear to face large degree of discrimination. They are severely affected by both monetary and non-monetary poverty; for example, they have low levels of literacy, are paid lower wages, and have less access to land or equal employment. According to a report based on a joint conference between the Cambodian Institute for Cooperation and Peace and the World Bank Institute, 43 percent of women are illiterate and 90 percent of these women are poor. According to the Cambodian Ministry of Women's and Veterans' Affairs, the following contribute to poverty—lack of opportunities for employment and education, lack of access to finance, mass illiteracy, food insecurity, malnutrition, human trafficking, powerlessness, no resources, overwork in wage employment and in households, discrimination in the labor market and in work places, and domestic violence. So many observers are surprised to learn that poverty rates are *not* higher among female-headed than male-headed

households in Cambodia. Likewise, female-headed households in neighboring Vietnam are no more likely to be in poverty than their male-headed counterparts.

8.2.2 Economic characteristics

Apart from income or consumption – which are typically used to define whether a household is poor – there are a number of other economic characteristics that correlate with poverty, most notably household employment and the property and other assets owned by the household.

There are several indicators for determining ***household employment***. Within this array of indicators, economists focus on whether individuals are employed; how many hours they work; whether they hold multiple jobs; and how often they change employment.

The ***property of a household*** includes its tangible goods (land, cultivated areas, livestock, agricultural equipment, machinery, buildings, household appliances and other durable goods) and its financial assets (liquid assets, savings and other financial assets). These indicators are of interest as they represent the household's inventory of wealth and therefore affect its income flow. Furthermore, certain households, especially in rural areas, can be poor in terms of income, but wealthy when their property is taken into consideration. Despite its importance, property is difficult to value in practice in any reliable way. First, one encounters the same problem of under-declaration. Second, it is very difficult to measure certain elements of property such as livestock. Finally, the depreciation of assets may be difficult to determine for at least two reasons: 1) the life span of any given asset is variable; 2) the acquisition of these assets occurs at different moments in each household. Therefore, property is more difficult to use than certain other elements in the characterization of poverty.

8.2.3 Social characteristics

Aside from the demographic and economic indicators, several social indicators are correlated with poverty and household living standards. The most widely used are measures of health, education and shelter.

Four types of indicators are normally used to characterize ***health*** in analyzing a household's living standards. These indicators include

- Nutritional status, for example, anthropometric indicators such as weight for age, height for age, and weight for height;

- Disease status, for example, infant and juvenile mortality and morbidity rates as related to certain diseases such as malaria, respiratory infections, diarrhea and sometimes poliomyelitis;
- The availability of health care services such as primary health-care centers, maternity facilities, hospitals and pharmacies, basic health care workers, nurses, midwives, doctors and traditional healers, and medical service such as vaccinations, access to medicines and medical information; and;
- The use of these services by poor and non-poor households.

Three types of indicators are normally used to characterize ***education*** in an analysis of household living standards. These include the level of education achieved by household members (basic literacy; years of education completed); the availability of educational services, such as proximity to primary and secondary schools; and the use of these services by the members of poor and non-poor households. For this last item, commonly used measures include children's registration in school, the dropout rate of children by age and gender and reasons for dropping out, the percentage of children who are older than the normal age for their level of education, and average spending on education per child registered.

Literacy and schooling are important indicators of the quality of life in their own right, as well as being key determinants of poor people's ability to take advantage of income-earning opportunities. Based on CSES data, Cambodia by 1993-94 had achieved a self-reported basic literacy rate of 67 percent among adults (older than age 15), implying a high degree of literacy among the poor. However, the literacy gap remained quite large, with literacy ranging from just over half of adults (58%) among the poorest 20 percent of the population to 77% among the richest 20% quintile. Much larger differentials appear in the distribution of schooling attainment: adults in the poorest quintile averaged 3.1 years of schooling, compared with 5.3 years among the richest quintile. Men averaged 5.1 years of education, compared with 3.2 years for women.

Shelter refers to the overall framework of personal life of the household. It is evaluated, by poor and non-poor household groups, according to three components (some of which overlap with the indicators mentioned above): housing, services, and the environment. The housing indicators include the type of building (size and type of materials), the means through which one has access to housing (renting or ownership), and household equipment. The service indicators focus on the availability and the use of drinking water, communications services, electricity, and other energy sources. Finally, the environmental indicators concern the level of sanitation, the degree of isolation (availability of roads and

paths which are usable at all times, length of time and availability of transportation to get to work) and the degree of personal safety.

> ***Example***. It is generally established that poor households live in more precarious, less sanitary environments, which contribute to the poorer health and lower productivity of household members. To illustrate: the data from the Cambodian Socio-Economic Survey of 1993/94 show that water and sanitation are especially important influences on health and nutritional status. The CSES showed that only 4 percent of the poorest quintile had access to piped water, while more than 17% of the richest quintile had the same. Similar differences are apparent in access to sanitation. Just 9% of the poor had access to a toilet in the home, while around half of the richest quintile did.
>
> Another indicator of housing standards is access to electricity. Here again the access of the poor lagged far behind. Access to electricity from a generator or line connection rose sharply with income, from a mere 1 percent among people in the bottom quintile to 37 percent of Cambodians in the richest quintile. Other indicators of household wealth include ownership of transportation. Access to bicycles is quite evenly distributed, with at least one half of households owning a bicycle in every quintile, even the poorest. However, access to cars, jeeps or motorbikes is very rare among the poor and rises sharply with income.

A summary of the main influences on poverty is provided in Table 8.1.

8.3 Analyzing the determinants of poverty: Regression techniques

Tabulated or graphical information on the characteristics of the poor is immensely helpful in painting a profile of poverty. However, it is not always enough when one wants to tease out the relative contributions of different influences on poverty. For example, data from the Vietnam Living Standards Survey of 1998 showed per capita expenditure to be significantly higher in female-headed households than in households headed by a man. However, after controlling for other influences – where the household lived, the size of the household, and so on – the effect proved to be statistically insignificant.

Table 8.1 Main determinants of poverty	
Regional characteristics	Isolation/remoteness, including less infrastructure and poorer access to markets and services Resource base, including land availability and quality. Weather (e.g. are typhoons or droughts common) and environmental conditions (e.g. frequency of earthquakes) Regional governance and management Inequality
Community characteristics	Infrastructure (e.g. is there piped water, access to a tarred road) Land distribution Access to public goods and services (e.g. proximity of schools, clinics) Social structure and social capital
Household characteristics	Size of household Dependency ratio (i.e. unemployed old and young relative to working age adults) Gender of head; or of household adults on average Assets (typically including land, tools and other means of production, housing, jewelry) Employment and income structure (i.e. proportion of adults employed; type of work – wage labor or self employment; remittance inflows) Health and education of household members on average
Individual characteristics	Age Education Employment status Health status Ethnicity

By far the most widespread technique used to identify the contributions of different variables to poverty is regression analysis. There are two main types of analysis:

(i) Attempts to explain the level of expenditure (or income) per capita – the dependent variable – as a function of a variety of variables (the "independent" or "explanatory" variables). The independent variables are typically of the type discussed above in section 8.2.

(ii) Attempts to explain whether a household is poor or not, using a logit or probit regression. In this case the independent variables are as in (i), but the dependent variable is binary, usually taking on a value of 1 if the family is poor and zero otherwise.

We now consider each of these in somewhat more detail.

A regression estimate shows how closely each independent variable is related to the dependent variable (e.g. consumption per capita), *holding all other influences constant*. There is scope for a wide variety of regressions; for instance the dependent variable could measure child nutrition, or morbidity, or schooling, or other measures of capabilities; the regressions could be used to examine the determinants of employment or labor income; or regressions could be used to estimate agricultural production functions (which relate production with information on type of crops grown per area, harvest, inputs into agricultural production, and input and output prices). For an accessible discussion and lots of examples, in the context of Vietnam, see Haughton et al., *Health and Wealth in Vietnam* (1999).

A typical multiple regression equation, as applied to poverty analysis, would look something like this:

$$\log(y_i / z) = \alpha_0 + \alpha_1 X_i^1 + \alpha_2 X_i^2 + \cdots + \alpha_n X_i^n$$

where z is the poverty line, y_i is (per capita) income or consumption, the X_j^n are the "explanatory" variables and the α_j are the coefficients that are to be estimated. Note that y_i/z is in log form, which is a common way of allowing for the log normality of the variable. Since we are interested in the determinants of individual poverty, but typically have information at the level of the household, it is standard (but in this context, not universal) to estimate the regression using weights that reflect the size of the household. The `regress` command in STATA is flexible and allows one to use weights.

The independent (i.e. right hand side) variables may be continuous variables, such as the age of the individual. But often we want to represent a categorical variable – the gender of the person, or the region in which he or she lives. In this case we need to create a "dummy" variable – for instance, the variable might be set to 1 if the person is a man and 0 for a woman. If there are, say, 10 regions in a country, each region would need to have its own dummy variable; but one of the regions needs to be left out of the regression, to serve as the point of reference.

Often we believe that the determinants of poverty differ from one area to the next, which would mean that there are differences in "structure." In this case we could estimate separate regressions for each area – for instance, for each region in a country. Sometimes it is sufficient to specify the regression equation in a way that is flexible enough to allow for such differences, for instance by allowing interactive effects: for example, one could create a variable that multiplies educational level by age, instead of estimating separate regressions for individuals in different age groups.

The fit of the equation is typically measured using $\bar{R}^2$ ("adjusted R squared"), which will vary between 0 (no fit) and 1 (perfect fit). There is no hard and fast rule for determining whether an equation fits well, although with household survey data one is often pleased to get an $\bar{R}^2$ of 0.5 or more.

We also need to know how much confidence to place in the accuracy of the coefficients as guides to the truth; this is commonly done by reporting t-statistics, which are obtained by dividing a coefficient by its standard error. The rule of thumb is that if the t-statistic is, in absolute terms, smaller than 2, then the coefficient is not statistically significantly different from zero (at about the 95% confidence level); in other words, we cannot be sure that we have picked up an effect, and it is possible that the coefficient just

reflects noise in the data. Many researchers prefer to report p-values, which give the confidence level directly; a p-value of, say, 0.03 indicates that we are 97% confident that the coefficient is not zero. So we hope to find low p-values (and we usually do when working with large data sets). Arbitrarily, it is standard to consider a coefficient to be statistically significant if the p-value is less than 0.05, but this rule is not graven in stone.

Table 8.2 shows typical regression output, from an example based on data from the Côte d'Ivoire. Here the dependent variable is the log of per capita household expenditure. The figures in parentheses are t-statistics. Separate regressions were estimated for households in urban, and in rural, areas, on the thinking that the determinants of poverty might be quite different in these two areas.

The results of the urban equation show that education is an important determinant of expenditure per capita. The coefficients for most of the educational variables are statistically significant and quite large; having an elementary education boosts income by approximately 38% relative to someone with no education; this comes from the coefficient of 0.38, and the fact that the dependent variable is in log form.

Table 8.2: Determinants of household spending levels in Côte d'Ivoire, c.1993

	Urban		Rural	
	Coefficient	**t-statistic**	**Coefficient**	**t-statistic**
Dependent variable: ln(expenditure/capita)				
Educational level of most-educated male				
Elementary	0.38	(5.3)	0.04	(0.6)
Junior secondary	0.62	(8.6)	0.08	(0.9)
Senior secondary	0.80	(9.6)	0.05	(0.4)
University	0.93	(9.4)	-	
Elementary	0.11	(1.7)	0.07	(1.0)
Junior secondary	0.24	(3.1)	0.27	(2.2)
Senior secondary	0.34	(4.1)	-	
University	0.52	(4.1)	-	
Value of selected household assets				
Home	0.06	(5.3)	-	
Business assets	0.04	(3.3)	0.16	(4.9)
Savings	0.08	(4.7)	-	
Hectares of agriculture land				
Cocoa trees			0.17	(4.3)
Coffee trees			0.04	(1.3)
Distance to nearest paved road			-0.04	(-2.9)
Distance to nearest market			-0.09	(-3.3)
Unskilled wage			0.37	(6.4)

Source: Adapted from Grosh and Munoz 1996, p.169, based on Glewwe 1990.

However, in rural areas education does not appear to explain expenditure per capita levels very much – a not uncommon finding. On the other hand, the infrastructure variables have substantial predictive power:

households located in villages that are nearer to both paved roads and public markets are better off, as are households living in areas with higher wage levels. The results raise further questions about the quality of education in rural areas and the importance of rural infrastructure in helping families grow out of poverty, which could be addressed in putting together a poverty reduction strategy.

It is vital to choose the independent variables carefully, and to be sure that they are truly exogenous. For instance, in the above example, one could have included income as an independent variable, along with education, assets, and the like. But that does not advance us much, because income is in turn determined by such variables as educational levels and household assets. In our drive to find the underlying causes of poverty, we need to dig deep to find variables that are indeed pre-determined. A good start is to work with the variables identified in Table 8.1.

When multiple cross-sectional surveys are available, the same regression can be repeated for different years to see how the association of certain correlates with income or consumption varies over time. Variations over time will be reflected in changes in coefficients or parameters. The results of repeated cross-section regressions can also be used to decompose variations in poverty in terms of changes in household characteristics, and changes in the returns to (or impact of) these characteristics (e.g., Baulch et al. 2001; van de Walle 2000; Wodon 2000). xxx

Some researchers prefer to use, on the left hand side, a binary variable that is set equal to 1 if the household is poor, and to 0 otherwise. Some of the information is lost by doing this, and the resulting logit or probit regression is relatively sensitive to specification errors, which is why this is rarely the preferred approach to take. However an analysis of this kind is likely to be useful when designing targeted interventions (e.g. educational vouchers for poor households), as it allows one to assess the predictive power of various explanatory variables used for means testing. It is also possible to undertake a multiple logit analysis, where the dependent variable could be in one of several categories – for instance, expenditure quintiles.

Recent research has explored more exotic forms of analysis, including non-parametric regression, classification and regression trees (CART models), and multiple-adaptive regression splines (MARS models). The goal of all such efforts is to unearth a parsimonious number of determinants of poverty, and quantify their effects, even when those effects are highly nonlinear.

References

Baulch et al. 2004 In Paul Glewwe, Nisha Agrawal and David Dollar (eds.), *Economic Growth, Poverty and Household Welfare in Vietnam*, World Bank 2004:

1. Ethnic Minority Development in Vietnam: A Socio-Economic Perspective [With Bob Baulch, Truong Thi Kim Chuyen and Dominique Haughton.]

Gibson. Cambodia. ??

VLSS

Haughton et al., 1999, *Health and Wealth in Vietnam.*

Grosh and Munoz. 1996.

van de Walle. 2000. Van de Walle, Dominique, and Dileni Gunewardena. 2001. "Sources of Ethnic Inequality in Vietnam." *Journal of Development Economics*, 65(1):177–207.

Wodon. 2000.

Chapter 9. Poverty Reduction Policies

Summary

Given a description and analysis of poverty, what policies may be invoked to reduce poverty?

There is a very strong link between economic growth and poverty reduction; Dollar and Kraay find, based on a study of 236 "episodes" worldwide that a 1% increase in per capita income is associated with a 1% increase in the incomes of the poor. The relationship is robust and has not changed over time. Greater economic openness,, the rule of law, and fiscal discipline all boost growth; democracy, and greater public spending on health and education, have no measurable effect on the incomes of the poor.

The World Bank classifies its anti-poverty activities into three groups:

a. **Fostering opportunity** – through well-functioning and internationally open markets, and investments in infrastructure and education.
b. **Facilitating empowerment,** which amounts to including people in the decision-making process. This requires government accountability, a strong media, local organizational capacity, and mechanisms for participation in making decisions.
c. **Addressing income security**, which tackles the problem of vulnerability. This calls for insurance programs, disaster relief procedures, and a sold public health infrastructure.

The chapter concludes with brief sketches of poverty reduction policies in India and Tanzania.

Learning Objectives

After completing the module on poverty lines, you should be able to:

1. Explain the methodology used by Dollar and Kraay to reach the conclusion that growth is good for the poor.
2. Evaluate the importance of other influences – including government spending, openness to trade, democracy, fiscal discipline, and the rule of law – on the growth of the incomes of the poor.
3. Describe what is meant by "pro-poor growth."
4. For each of the three groups of anti-poverty activities identified by the World Bank, i.e.
   - promoting opportunity
   - facilitating empowerment, and
   - enhancing income security

   justify the importance of each broad activity and identify specific policies within each of these activities that are likely to work to reduce poverty.

Previous chapters have discussed the concept of poverty and well-being, the various indicators used to measure poverty, the idea of poverty profiles, and the factors that determine poverty. In this section, we consider poverty reduction strategies, as well as touch upon certain policy implications.

9.1 Is growth good for the poor?

Few economists doubt that economic growth is necessary for the long-term reduction of poverty. But how close is the link between the two?

If the incomes of the poor rise closely in line with incomes overall, then the key to poverty reduction is rapid economic growth; on the other hand if the relationship is weak, then other policies, such as targeted subsidies, are likely to be important and the concept of "pro-poor growth" might have some relevance.

David Dollar and Aart Kraay address the problem directly, in a paper entitled "Growth is Good for the Poor" that appeared in March 2000 (see www.worldbank.org/research). They gathered information on the per capita income of the poor (the bottom quintile of the income distribution) and on overall per capita income. The data come from 80 countries over four decades, and Dollar and Kraay were able to piece together 236 "episodes," - periods during which it was possible to measure changes in the income of the poor and of the country overall.

They first regressed the log of per capita income of the poor (ln(poor)) on overall per capita income (ln(inc)) and got

$$\ln(\text{poor}) = 1.07 \ln(\text{inc}) - 1.78. \quad R^2=0.87. \qquad (1)$$

This relationship, and the underlying data, are reproduced in Figure 9.1 (Figure 1a from Dollar and Kraay) below. Two points are worth noting: First, the relatively high value of R^2, which says that 87% of the variation in the log of per capita income of the poor is associated with changes in the log of per capita income overall. Second, the coefficient on the ln(inc) term is 1.07, which means that if average incomes were to rise by 10%, the incomes of the poor would rise by about 10.7%. This coefficient is not significantly different from 1, so perhaps it would be wiser to conclude that the incomes of the poor tend to rise *in line with* those of the country as a whole.

As an alternative, Dollar and Kraay regressed the *change* in ln(poor) on the change in ln(inc). In this case they found (see Figure 9.2, which reproduces Figure 1b from Dollar and Kraay)

$$\Delta \ln(\text{poor}) = 1.17 \ln(\text{inc}) - 0.00. \qquad R^2=0.52.$$

The fit is weaker in this equation, with only about half of the variation in the change in the log of incomes of the poorest quintile being attributable to changes in the log of overall income. The elasticity (i.e. 1.17) is still close to unity.

In order to test the robustness of this finding, Dollar and Kraay estimate a number of variations on the original equation - for different time periods, for countries that are growing and countries that are shrinking, for low- and high-growth countries, for poor countries and rich. The details are summarized in their Figure 2 (reproduced below), and show that the elasticity of poor incomes with respect to overall incomes is close to 1 in every case. They also address the issues of measurement error (it washes out), omitted variable bias (the use of instruments gives similar results) and endogeneity (systems estimators also give similar results). And when they include additional variables, the key elasticity remains close to 1, as Table 9.1 (their Table 5, reproduced below) shows.

There are several important conclusions:

1. The per capita income of the poorest quintile grew in line with overall per capita GDP for the 80 countries in the same over the four decades covered. In short, growth matters.
2. The relationship in equation (1) has not changed over time, and applies both in rich and poor countries. In short, growth still matters.
3. The incomes of the poor do *not* fall disproportionately during an economic crisis.
4. Greater economic openness, the rule of law, fiscal discipline, and low inflation all contribute to (or are at least associated with) faster economic growth, and in this manner help the poor.
5. Democracy, and higher public spending on health and education, do not have a measurable effect, one way or another, on the incomes of the poor.

Figure 1: Growth and the Poor

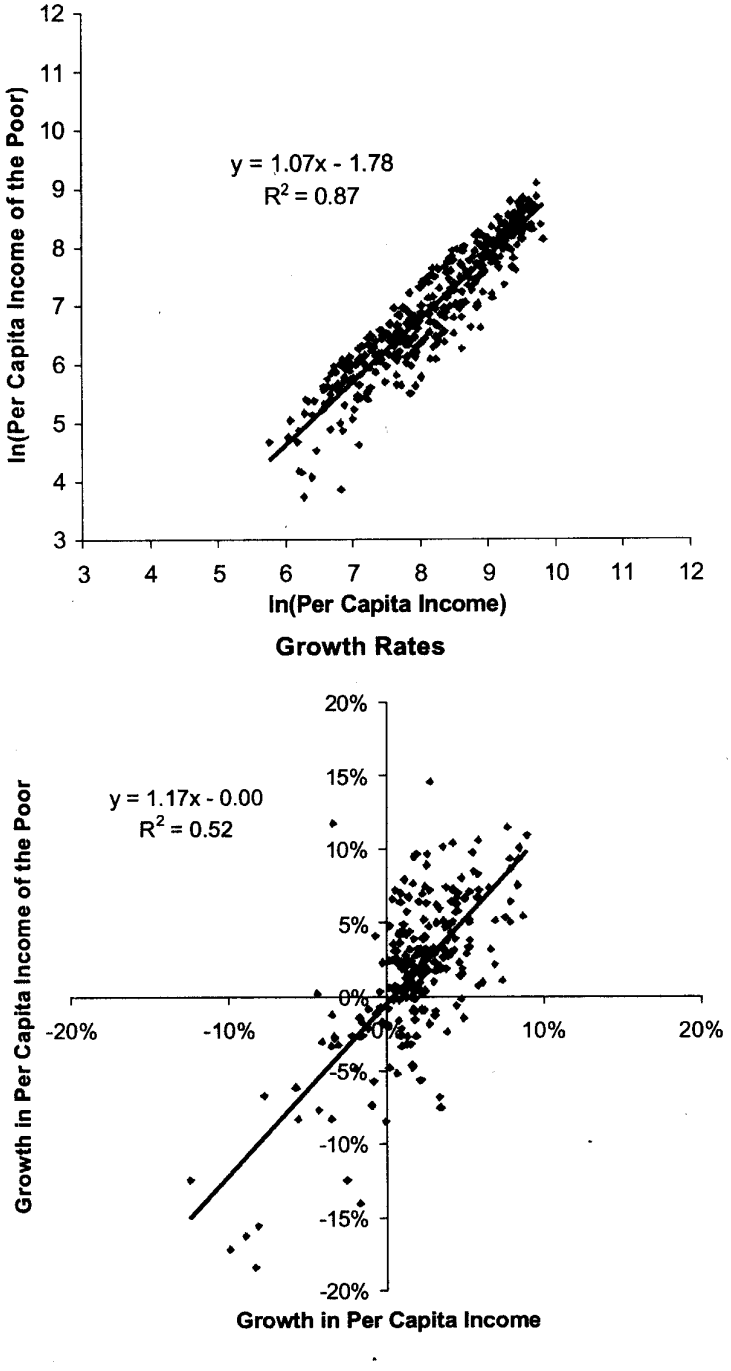

Figure 2: Variants on Basic Regression

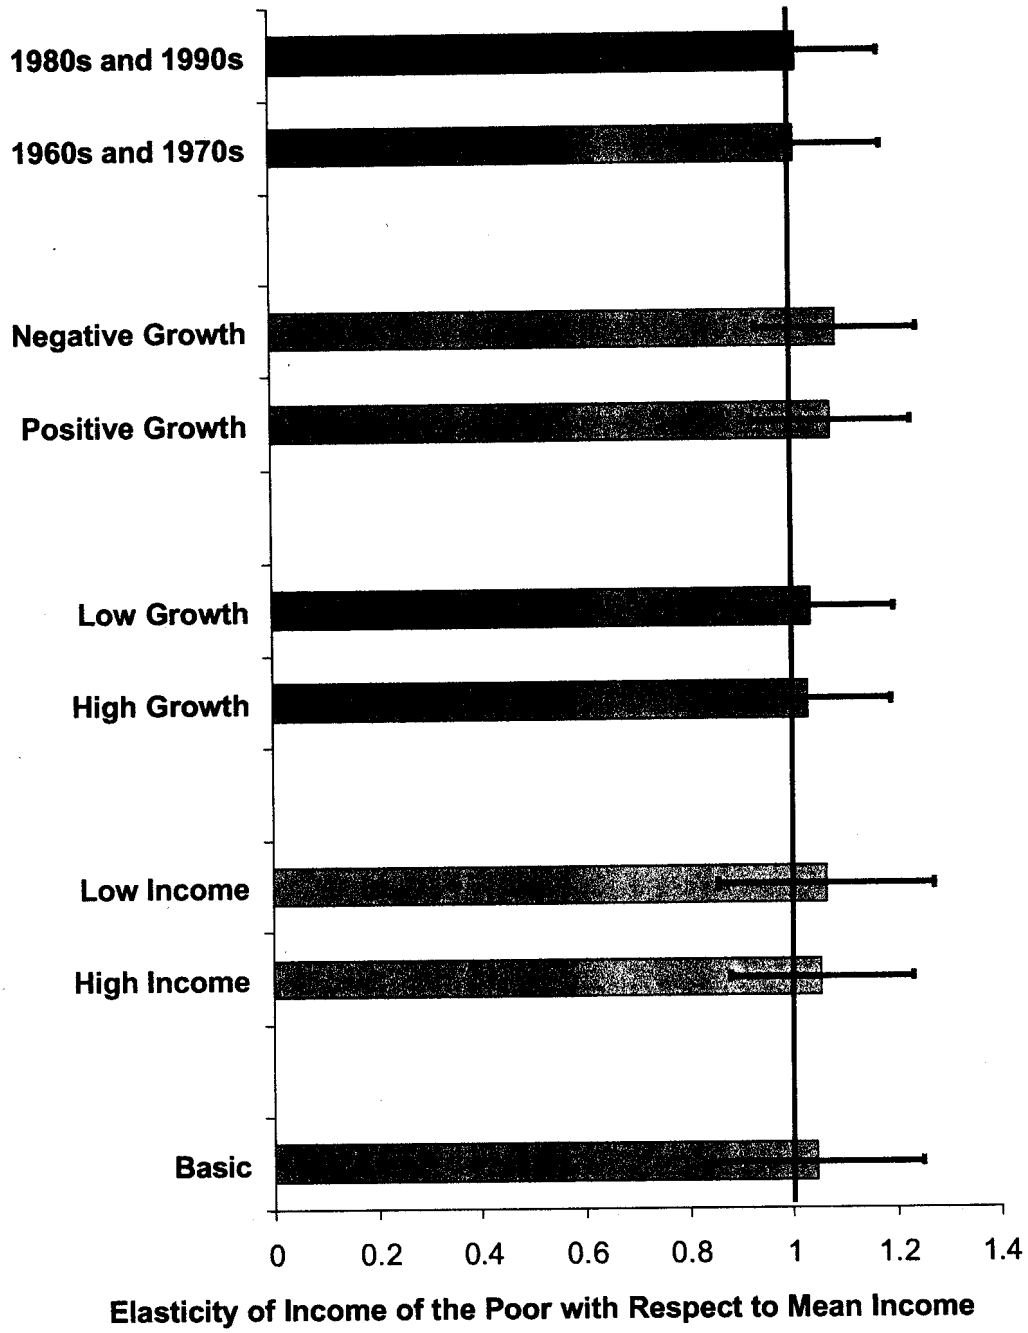

Table 5: Pro-Growth Policies and the Poor

(Dependent Variable is ln(Per Capita Income of the Poor))

	(1)	(2)	(3)	(4)	(5)
ln(Per Capita Income)	1.055	1.063	1.021	1.042	1.104
	(0.102)	(0.088)	(.080)	(.092)	(.094)
(Exports+Imports)/GDP	0.004				-0.004
	(0.055)				(.013)
ln(1+Inflation)		-0.134			-0.021
		(0.089)			(.012)
Government Consumption/GDP			0.0001		0.0005
			(.0001)		(.001)
Rule of Law				0.005	-0.041
				(.067)	(.065)
# of Observations	213	232	214	235	210
P-OID	0.166	0.384	0.208	0.271	0.499

Notes: This table reports the results of adding the indicated control variables to the system estimator of Equations (2) and (2'). The row labelled P-OID reports the P-value associated with the test of overidentifying restrictions. Standard errors are corrected for heteroskedasticity and for the first-order autocorrelation induced by first differencing using a standard Newey-West procedure.

Other researchers have also found that poverty trends tracked growth trends very closely in 1980s and 1990s. According to Chen and Ravallion (2000), on average, growth in the consumption of poorest fifth of the population tracked economic growth one-for-one over this period. In the vast majority of countries that they study, growth led to rising consumption in the poorest fifth of the population, while economic decline led to falling consumption. So pro-growth policies in most cases actually are also pro-poor.

9.2 Pro-poor growth

If the incomes of the poor are closely tied to overall economic growth, how much room remains for a poverty reduction policy per se? Put another way, is there any substance in calls for "pro-poor" growth?

In a recent controversial paper, Aart Kraay (2004) argues that "in the medium run, most of the variation in changes in poverty is due to growth, suggesting that policies and institutions that promote broad-based

growth should be central to pro-poor growth." He goes on to argue, "most of the remainder is due to poverty-reducing patterns of growth in relative incomes," but "cross-country evidence provides little guidance on policies and institutions that promote these other sources of pro-poor growth." In other words, we do not know enough about what drives pro-poor growth – roughly, growth accompanied by a reduction in inequality – to be in a position to design viable pro-poor policies.

That has not stopped the World Bank and others from trying! The World Bank (2000) separates its anti-poverty – as distinct from pro-growth – activities into three groups: promoting opportunity, facilitating empowerment, and enhancing (income) security. We now consider each in some more detail.

9.2.1 Opportunity

As argued in chapter 8, the lack of material opportunities such as jobs, credit, and public services, including schools and health services, is a direct cause of poverty.

Well-functioning markets are important in generating sustainable growth and expanding opportunity for poor people because poor people rely on formal and informal markets to sell their labor and products, to finance their investments, and to insure against risks. For example, recent studies have examined the impact of market-friendly policies – such as openness to international trade, low inflation, a moderate-size government, and strong rule of law – on the incomes of poor people in a large cross-country sample. The findings: these policies on average benefit poor people more than others. Case studies of Chile, China, Ghana, Uganda and Vietnam show that agricultural reforms have helped raise producer prices for small farmers by eliminating marketing boards, changing real exchange rates through broader economic reforms, lowering tariffs and eliminating quotas (e.g. Haughton and Kinh 2003).

The human, physical, natural, financial and social assets that poor people possess – or have access to – affect their prospects for escaping poverty because these assets can enable poor people to take advantage of opportunities. Expanding the assets of poor people can strengthen their position and their control over their lives. A recent study of irrigation in Vietnam (van de Walle 2000a) has uncovered important complementarities between education and gains from irrigation. Households with higher education levels received higher returns to irrigation.

Policies to Enhance Opportunity

The World Bank argues that creating more opportunities involves complementary actions to stimulate overall growth, make markets work for poor people, and build their assets, including addressing inequalities in the distribution of endowments such as education. More specifically:

- Encouraging effective private investment is essential because investment and technological innovation are the main drivers of growth in jobs and labor incomes. Fostering private investment requires reducing risk for private investors—through stable fiscal and monetary policy, stable investment regimes, sound financial systems, and a transparent business environment. Certainly, the rule of law and anti-corruption measures are also important. Special measures are frequently required to ensure that micro enterprises and small businesses can participate effectively in markets that are more vulnerable, yet employ a large number of poor people. For example, ensuring access to credit, lowering transaction costs of reaching export markets, and reducing restrictions on the informal sector will all help creating a sound business environment for poor households and small firms. Public investment in expanding infrastructure and communications, and upgrading the skills of the labor force have to complement private investment to enhance competitiveness and create new market opportunities.
- Opening to international markets offers an important opportunity for income growth as long as countries have the infrastructure and institutions to stimulate a strong supply response (e.g. call centers in Ghana, garment factories in Vietnam). Therefore, the opening needs to be well designed with special attention to bottlenecks.
- Building human, physical, natural and financial assets that poor people own or can use requires actions on at least two fronts. First, sufficient public spending on basic social and economic services such as primary education. Second, reform public delivery or privatize those services in order to ensure good quality service delivery.
- Addressing asset inequalities across gender, ethnic, racial and social divides. Special action is required to tackle socially-based inequality such as concentrated farm land ownership in rural communities, under-schooling of girls relative to boys, and the limited independence of women due to lack of access to productive means. Ethnic inequalities can easily flare up into violence, which in turn can set back economic development for a generation.
- Getting infrastructure and knowledge to poor areas. Special action is also needed in order to improve the social and economic infrastructure in poor and remote areas, which to a great extent also contribute the poverty problem. Similarly, basic urban services should be provided to city slums so that urban poor people may have chance to participate more actively in over growth.

9.2.2 Empowerment[8]

The premise underlying an emphasis on empowerment is that a lack of representation in the process of policy-making, due to social and institutional barriers, has impeded poor people's access to market opportunities and to public sector services. It follows that empowerment – defined succinctly as "including people, who were previously excluded, in the decision making process" – should help. Unfortunately, there is very little empirical evidence, to date, on how well empowerment policies along the lines discussed below contribute to reducing poverty.

Broadly, empowerment refers to being able to make informed decisions and choices effectively. But there is some disagreement about the true content of empowerment. Mahatma Gandhi emphasized self-reliance; Paolo Freire stressed the need for conscientization, for helping the poor to learn about and perceive "social, political and economic contradictions" and then to stir to act against "the oppressive elements of society." E.F. Schumacher, author of *Small is Beautiful*, argues that empowerment follows when one makes up deficiencies in education, organization and discipline. The World Bank finesses these differences by defining empowerment as "the expansion of assets and capabilities of poor people to participate in, negotiate with, influence, control, and hold accountable institutions that affect their lives." It sees the four major elements of empowerment as (i) access to information, (ii) inclusion/participation, (iii) accountability, and (iv) local organizational capacity.

State institutions must be responsive and ***accountable*** to poor people. In nearly every country the public sector often pursues activities that are biased against poor people, and poor people have trouble getting prompt, efficient service from the public administration. Accountability is helped when there is good ***access to information***.

> ***Example:*** The Public Expenditure Tracking Survey (PETS) conducted in 1996 in Uganda found that only 22% of the central government funds intended to support schools run at the local level were reaching their intended destination. By 1999-2000, after the government made the budgetary transfers public via the media and required schools to share financial information, 80-90% of the funds began to reach the schools for which they were intended.

[8] The material in this section is derived from a survey paper entitled *Employment*, by Zeynep Orhun.

Amartya Sen sees poverty as consisting of a "deprivation of capabilities," so that the poor have inadequate resources (financial, informational, etc.) to participate fully in society; in short, they are socially excluded. It follows that ***inclusion***, which encompasses economic and political participation, is inherently part of the solution to poverty. The process of including the poor is likely to require the development of their social capital, the "features of social organization, such as trust, norms, and networks, that can improve the efficiency of society by facilitating coordinated actions." Social capital takes time to build, but contributes to stronger ***local organizational capacity.***

Good social institutions — kinship, community organizations, and informal networks—can play an important role in poverty reduction. For example, many development programs succeed because they mobilize local groups of project beneficiaries in program design and implementation. On the other hand, when social institutions are week, fissures such as ethnic cleavages can explode into open conflict; most of the world's poorest countries have experienced civil war within the past generation.

Some social norms and practices help generate and perpetuate poverty. Discrimination on the basis of gender, ethnicity, race, religion, or social status can lead to social exclusion and create barriers to upward mobility, constraining people's ability to participate in economic opportunities and to benefit from and contribute to economic growth. For example, one cross-country study (Klasen, 1999) indicates that countries that invest in girls' education have higher rates of economic growth.

It is difficult to empower the poor if decision-making is concentrated in a far-away capital city. Hence the conclusion that a major component of empowering the poor is the need to decentralize power, particularly through delegating it to sub-national levels of government, and privatizing some activities (e.g. grain marketing). Decentralization is not, however, a panacea; when done badly, power may be captured by local elites, who may be even less concerned about the poor than the central government. In India, for instance, the state of Kerala has used its powers to spread development widely, while in the state of Bihar local decision-making has not been particularly beneficial to the poor.

Empowerment is difficult to measure. The UNDP's ***Gender Empowerment Measure (GEM)*** includes indicators such as male and female shares of parliamentary seats, managerial positions, and earned income, but also has serious limitations in that it does not include information on the informal sector, or on such items as the right to vote. By design, the GEM focuses on gender empowerment, and not on the empowerment of the poor per se.

Policies to enhance empowerment

To empower poor people, policies need to facilitate active collaboration among poor and other groups in society include strengthening the participation of poor people in political processes and local decision-making; making changes in governance that make public administration, legal institutions, and public service delivery more efficient and accountable to all citizens; and removing the social barriers that result from distinctions of gender, ethnicity, race and social status. Worthy as this sounds, it is not at all obvious how to achieve such changes, but here are some policies that have been suggested:

- To improve access to ***information***, encourage the development of the media. For instance, Besley and Burgess show that there is a robust link between media development and government responsiveness in India; states with higher newspaper circulation also undertake more extensive relief efforts in the wake of natural disasters.
- To increase ***participation and inclusion***, it helps to institutionalize transparent, democratic and participatory mechanisms for making decisions and monitoring implementation. In this context, it may also be useful to provide legal assistance to poor people who usually have limited access to the legal system.
- ***Accountability*** is increased by strengthening the mechanisms used to monitor the performance of public administrations and by providing access to budgetary information and participatory mechanisms. There are many possible ways to do this, including:
    - Publish complete and timely budgetary information. Vietnam does not do this, for instance, so it is impossible to hold the government to account for how it spends its money.
    - Institutional and Governance Reviews (IGR), which use surveys and other quantitative measures to analyze the functioning of public institutions.
    - Citizen Report Cards, which allow citizens to express their opinions on the performance and quality of government services.
    - World Bank Corruption Surveys, which are designed to extract information on corruption from households, the private sector, and public officials. Based on such a survey, for instance, Albania requested an anti-corruption program to undermine patronage in judicial and civil service appointments.
    - Public Expenditure Tracking Surveys (PETS), which have helped ensure that budgeted funds get to their intended recipients in places such as Ghana and Uganda.
    - Private Enterprise Surveys of the Business Environment, and Investor Roadmaps. These indicate the problems and costs faced by entrepreneurs.

- Participatory Poverty Assessments (PPA). Using focus groups, in-depth interviews and other measures, these complement survey data to help build a more detailed picture of the nature and roots of poverty; they have been influential in Vietnam, for instance.

- To increase ***local organizational capacity,*** it helps to
  - Promote decentralization and community development to enhance the control that poor people and their communities have over the services to which they are entitled. Decentralization needs to be combined with effective participation and monitoring mechanisms.
  - Promote gender equality by promoting women's representation in decision-making and providing special assistance for women's productive activities.
  - Tackle social structures and institutions that are obstacles to the upward mobility of poor people by fostering debate over exclusionary practices and supporting the socially excluded participating into political processes.
  - Support poor people's social capital by assisting networks of poor people to engage with market and non-market institutions so as to strengthen their influence over policy.

9.2.3 Income security

Poor people are exposed to a wide array of risks that make them vulnerable to income shocks and losses of well-being. Reducing poor people's ***vulnerability*** to ill health, economic shocks, natural disasters, and violence enhances well-being on its own and encourages investment in human capital and in higher-risk, higher-return activities as well.

Households and communities respond to their risk exposure through diversification of assets and sources of income and various types of self-insurance and networks of mutual insurance mechanisms. For instance, some family members may travel to the cities to seek work, sending remittances home; if they cannot find work, then return home. Or farmers may store grain from one season to the next, in case the crops fail. In some very poor countries, such as Mali, some very poor rural women wear large gold ornaments – in effect carrying their savings, which could be sold if necessary to tide the household over a bad year.

Mechanisms such as these help to reduce the risk or soften the impact, but the effect may be limited. To counter the incentive and information problems that exclude poor people from many market-based insurance mechanisms, the state has a special role in providing or regulating insurance and setting up

safety nets. Health, environmental, labor market, and macroeconomic policies can all reduce and mitigate risk.

Large adverse shocks – economic crises and natural disasters – cause poor people to suffer not only in the short run. They undercut the ability of the poor to move out of poverty in the long run as well, by depleting their human and physical assets, which may be irreversible. So it is crucial to prevent economic crises and natural disasters, as well to protect poor people when they occur.

Policies to enhance security and reduce vulnerability

National programs to manage economy-wide shocks and effective mechanisms to reduce the risks faced by poor people, as well as helping them cope with adverse shocks when they occur, are useful. Appropriate measures might include:

- Formulating programs to helping poor people manage risk. Microinsurance programs, public works programs and food transfer programs may be mixed with other mechanisms to deliver more effective risk management.
- Developing national programs to prevent and respond to macro shocks—financial or natural.
- Supporting minority rights and providing the institutional basis for peaceful conflict resolution, to help prevent civil conflict and mobilize more resource into productive activities.
- Tackling health problems including widespread illnesses such as malaria and tuberculosis, as well as moderately common but serious conditions such as HIV/AIDS.

There is no simple, universal blueprint for implementing this strategy. Each developing country needs to prepare its own mix of policies to reduce poverty, reflecting national priorities and local realities. Given the important complementarities among these three dimensions, an effective poverty reduction strategy will require action on all three fronts, by all stakeholders in society—government, civil society, the private sector and poor people themselves.

9.3 Practice and good examples

Any good poverty reduction plan has to be based on a comprehensive poverty analysis that identifies the nature and evolution of poverty, the profile of poor people, and all contributing factors of poverty. Building on an accurate understanding of poverty, the strategy for poverty reduction has to prioritize the

poverty reduction goals and take into account complementarities and compatibilities of various policy tools. Then specific implementation modules including resource allocation and monitoring mechanisms have to be designed.. The participation of poor people at the last two stages is essential since overall they are the main actors in the fight against poverty. We finish this chapter with brief sketches of two interesting examples of poverty reduction strategies.

9.3.1 India

India suffers severe deprivations in education and health—especially in the North, where caste, class, and gender inequities are particularly strong. In studies in Bihar and Uttar Pradesh, poor women and men emphasized their extreme vulnerability and the ineffectiveness of state institutions, from schools to police.

In the past, poverty reduction in India lagged behind that of East Asia because of slower growth and significantly less progress in promoting mass education and basic health. More recently, however, growth has accelerated and there is a growing consensus that poverty has fallen substantially.

There are also marked differences within India—with the South, particularly the state of Kerala, having sharply better education and health. Kerala has life expectancies greater than those in Washington, D.C., despite vastly lower income levels. The effectiveness of public action in Kerala has been attributed to its strong tradition of political and social mobilization.

What are the priorities for action in India? Accelerated poverty reduction requires continued solid economic growth, which in turn demands further liberalization, especially in agriculture, and better provision of infrastructure, sorely lacking in most of India. In areas with deep deprivation in health and education, the development of social infrastructure is critical. Expanding education and health services requires that state governments reverse the deterioration in their fiscal positions, as subsidies to the loss-making power sector crowd out spending in the social sectors. The higher spending needs to be matched by better service provision. This requires deep improvements in governance, often weakest in India's poorest regions, and in combating teacher absenteeism. Also needed is more equitable service provision, which requires empowering women and members of lower castes.

9.3.2 Tanzania

Since independence, the government has been focused on three development problems: ignorance, disease and poverty. National efforts to tackle these problems were initially channeled through centrally directed, medium-tem and long-term development plans; despite high levels of foreign aid, these efforts were a complete failure, and poverty was higher in 1990 than at the time of independence. Despite sustained effort over the past decade, and an acceleration in economic growth, half of the population today is considered basically poor, and approximately one third lives in extreme poverty.

The current poverty reduction strategy takes a decentralized approach, based on broad consultation with all stakeholders. The Zonal Workshops, which included a large number of representatives of the poor at the village level, identified priority concerns and requirements for poverty reduction. Through an assessment of the poverty profiles and trends using 6 household surveys of various purposes during the past two decades, the key findings include: a very high incidence of poverty (48% in 1991/92); poverty is more widespread in rural areas than in urban centers; the poor are concentrated in subsistence agriculture; urban poverty is also acute in urban areas other than Dar es Salaam, and poor people are mostly in the informal sector; youth, the old and large households are more likely to be poor; the incidence of poverty declined between 1983-1991 and 1991-1993, but rose during 1993-1998; infant mortality rate remains to be high and AIDS is the leading killer of youths in Dar es Salaam and several other sampled districts; malnutrition continues to be severe; access to clean drinking water is still limited and the majority of poor people have no access to piped water and, as a result, have much higher exposure to cholera and waterborne diseases; the judiciary system has corruption problems.

Two participatory poverty assessment studies have suggested the following causes of poverty and obstacles to development: insecure land tenure, lack of access to agricultural inputs, credit, technology, transportation, markets, and quality health services, vulnerability to shocks, gender inequality. These factors were reiterated in the Zonal Workshops. The information then was summarized for each dimension of poverty and further discussed with other stakeholders.

The poverty reduction strategy aims at (1) reducing income poverty; (2) improving human capabilities, survival and social well-being; and (3) containing the extreme vulnerability among the poor. The key instrument is high and equitable growth. Specifically the strategy aimed (optimistically) to raise, within three years, the annual GDP growth rate from 5.2% to 6%. Agriculture was expected to increase by 5% annually on average, compared to 3.6% during 1990-1998. The growth of the industrial and service

sectors was expected to be higher than that of agriculture. At these growth rates, the incidence of poverty is expected to drop to 42% by 2003 and 24% by 2010. In pursuit of the above poverty reduction objectives, reforms aimed at bolstering market efficiency and factor productivity will be implemented. Besides maintaining low inflation and a small fiscal imbalance, special budgetary attention is to be given to rural, export, and private sectors development programs.

Specific goals were also identified for other dimensions of poverty such as education and health, and especially HIV/AIDS related issues. Twenty seven indicators for the monitoring and evaluation of the poverty reduction strategy have been selected and at least two observations will be made during the next three years. These key welfare indicators will be at the heart of further discussion and assessment of the poverty reduction strategy.

References

Besley and Burgess.
Chen and Martin Ravallion. 20000.
Dollar, David and Aart Kraay. 2000. "Growth is good for the poor," Available at www.worldbank.org/research .
Friere, Paolo.
Haughton, Jonathan and Hoang Van Kinh. 2004.
Klasen. 1999.
Kraay, Aart. 2004.
Schumacher, E.F. *Small is Beautiful*
Sen, Amartya.
van d Walle. 2000. [Irrigation in Vietnam]
Uganda. Public Expenditure Tracking Survey.
United Nations Development Program. [Gender Empowerment Measure]
World Bank. 2000.

Chapter 10. International Poverty Comparisons[9]

Summary

The central target of the Millennium Development Goals (MDG) is to half, between 1990 and 2015, the proportion of people worldwide whose income is less than a dollar a day. To measure progress towards this goal it is necessary to compare poverty rates across countries.

This requires an absolute measure of poverty, typically in terms of income or expenditure per capita. The World Bank reports that the dollar-a-day poverty rate fell from 40% of the population of LDCs in 1981 to 21% by 2001, with the biggest fall in East Asia (from 58% to 15%) and no reduction in Africa.

These figures are sensitive to what exchange rate is used; strictly, it should reflect purchasing power parities (PPP) and apply to a basket of goods and services consumed by the poor. In practice, PPP exchange rates are often based on econometric inference rather than measured directly, and are rarely applied to poor-person baskets of goods and services. Small changes in assumptions show Uganda's poverty rate to be either 85% of 27%!

Household survey data understate income (and expenditure). When reconciling the results with national accounts it is tempting, but often misleading, to gross up the income of every household by same proportion in order to achieve consistency with the measure of national income. In some cases this appears to understate poverty.

Over the long run, economic growth powers poverty reduction. But in the short-run the link is weaker, and so even if the world economy is growing, one cannot assume that the poverty rate is necessarily falling.

Learning Objectives

After completing the module on poverty lines, you should be able to:

1. Describe the main target of the Millennium Development Goals.
2. Justify the need to make international comparisons of poverty.
3. Identify those parts of the world where poverty has fallen most quickly, and least quickly, since 1981.
4. Explain why international poverty comparisons require the use of purchasing power parity (PPP) exchange rates.
5. Explain how PPP exchange rates should be computed for international poverty comparisons, and how they are typically computed in practice.
6. Explain what basket of goods and services should be used for international poverty comparisons.
7. Illustrate, using the example of Uganda or another country, how the methodological choices can make a large difference to measured poverty rates.
8. Summarize the challenges involved in reconciling household survey data (where income and expenditure are typically undervalued) with national accounts data.
9. Assess whether world poverty is falling, and if so, how quickly.
10. Recognize that while economic growth drives poverty reduction in the long run, this need not be the case in the short run.

[9] This chapter is mainly based on Kevin Carey, *Poverty Analysis for Macroeconomists,* World Bank Institute, August 2004.

10.1 Introduction

The first target of the Millennium Development Goals (MDG) is to halve, between 1990 and 2015, the proportion of people worldwide whose income is less than one dollar a day.[10] This naturally leads to a simple question: are we on track to meet this goal? But to answer this question we need to be able to compare poverty rates across countries.

The World Bank and other donor and lender agencies have limited resources. Many are interested in channeling these scarce resources to countries where poverty is especially high. But to do this, one again need to be able to compare poverty rates across countries.

Cross-country comparisons of poverty face two problems:

a. What exchange rate to use in order to apply a common poverty line worldwide; and
b. How to reconcile survey data with national accounts. The difficulty here is that when one adds up consumption based on household budget survey (HBS) data, the result is typically smaller than one would expect based on national income data.

These problems are technical but they are not trivial. Depending on the exchange rate used, the percentage of Ugandans living on less than a dollar a day is 85% or 27%! And while Bhalla (2002) argues, using national accounts data, that the Millennium Development Goal for poverty reduction is close to being met, others strongly disagree.

In this chapter we examine these issues in more detail.

10.2 Overview of poverty analysis

As described in earlier chapters, the key steps in the measurement of poverty are to: specify a minimal socially acceptable level of income or consumption (the *poverty line*), implement a representative survey in which the corresponding income or consumption concept is measured, and choose and calculate a specific poverty measure. The most common implementation of these steps is to have a fixed monetary consumption-based threshold for poverty, with data coming from a household

[10] Information on the overall MDGs and progress towards them can be found at http://www.developmentgoals.org/.

survey, and poverty measured as percentage of individuals with per capita consumption below the poverty line (the *headcount* measure).

Even at this broad level, notice the subtle restrictions that have already emerged; we are defining poverty in absolute and not relative terms; we tend not to focus on non-monetary measures of well-being, such as health; poverty is a concept that applies to individuals but is measured from household data; and in practice we nearly always use the headcount measure, even though this is just one of many possible measures.

Chapter 2 has discussed the need to identify the preferred indicator of welfare in terms of which the poverty line will be specified. For economists, the choice of indicator boils down to income versus consumption. There is a theoretical preference to measure consumption poverty, especially for developing countries where labor market participation (and the associated income paper trail) can be extremely erratic. First, it is consumption that appears in utility functions. Second, consumption corresponds to "permanent income." Third, the conceptual advantage of consumption over income is strengthened by data considerations. Measurement of income suffers from deliberate understatement, measurement error, and omission of key components (e.g. capital gains on infrequently marketed assets).

But consumption also poses difficult measurement issues, especially bearing in mind that it requires data on both quantities and prices. There is relatively good experience worldwide with measurement of non-durable consumption. But we should also be including the service flow from all durable goods, and household surveys attempt to do this. With a perfect rental market in durable goods, this would be easy: consumption service flow would correspond to the market or shadow rent on the durable good, which in turn would equal depreciation plus opportunity cost. But durable goods markets exist for very few goods and can be very thin even when they do exist. This forces arbitrariness in concerning assumptions for depreciation and opportunity cost for the poor; in particular, the standard usage of market interest rates as a measure of opportunity cost may make little sense for the poor with constrained access to capital markets.

10.3 International poverty comparisons

Let us briefly examine the regional aggregate data on poverty outcomes which form the basis for the global dialog on the poverty MDG. The World Bank's estimates of the proportion of the population

living below the dollar-a-day level is shown immediately below (from www.worldbank.org/research/povmonitor, accessed August 8, 2004). Of interest is the massive reduction in the poverty rate in East Asia between 1981 and 1987, largely due to the drop in poverty in China; the rise in poverty in Europe and Central Asia (mainly in the states of the Former Soviet Union); and the upward drift of poverty rates in sub-Saharan Africa.

Headcount Indicies: Percentage of Population Living Below $1 Per Day

Click on the region below to access detailed country information	Poverty Rate (% below $1.08/day at 1993 PPP)								For Data Source
	1981	1984	1987	1990	1993	1996	1999	2001	
East Asia and Pacific	57.7	38.9	28.0	29.6	24.9	16.6	15.7	14.9	click here
China	63.8	41.0	28.5	33.0	28.4	17.4	17.8	16.6	
Europe and Central Asia	0.7	0.5	0.4	0.5	3.7	4.2	6.3	3.7	click here
Latin America and the Caribbean	9.7	11.8	10.9	11.3	11.3	10.7	10.5	9.5	click here
Middle East and North Africa	5.1	3.8	3.2	2.3	1.6	2.0	2.6	2.4	click here
South Asia	51.5	46.8	45.0	41.3	40.1	36.6	32.2	31.3	click here
India	54.4	49.8	46.3	42.1	42.3	42.2	35.3	34.7	
sub-Saharan Africa	41.6	46.3	46.8	44.6	44.0	45.6	45.7	46.9	click here
Total	40.4	32.8	28.4	27.9	26.3	22.8	22.2	21.1	

In international comparisons it is common to try to measure the percentage of the population living on less than one US dollar per person per day, as above, or the two dollars per day count that is more relevant for the lower middle income countries. The dollar-a-day standard refers to the purchasing power of a "dollar", actually $1.08, in 1993 (equivalent to $1.41 in 2004 prices). Chen and Ravallion (2001, 2004) argue that this poverty line is representative of the actual poverty line in very poor countries. The usual procedure is to convert a 1993 dollar to local currency and then update to the year of interest using the local Consumer Price Index (CPI) to get the line in terms of prices at the desired survey date.

> *Example:* $1 was worth 25 Thai baht in 1993. Between 1993 and 2004, consumer prices in Thailand rose by 47%. Thus the poverty line in Thailand would now (mid-2004) be 36.7 baht per person per day. It is worth noting that the exchange rate in mid-2004 was 41.5 baht per person per day.

The first problem is determining the purchasing power of a dollar in each country. Clearly $1 buys less in the US than in India (where it bought 46 rupees in August 2004). Someone living on $500 per month in the US would be poor; in India they would be comfortably off.

So why do exchange rates not reflect the relative purchasing powers of different currencies? The answer is that while tradable goods (e.g. TVs, basmati rice) have similar prices everywhere – allowing for transport costs of course – this is not true of non-tradable goods. For instance, a simple haircut in Hanoi (Vietnam) costs $0.33, while in Boston (US) it costs $12. Despite this price differential, it does not make sense for people to fly from Boston to Hanoi to get a haircut!

The standard solution to the "exchange rate problem" is to recompute incomes, for different countries, in a common set of "international" prices. First done on a large scale by the UN-sponsored International Comparison Project, this is the basis for purchasing power parity (PPP) cross-country comparisons of per capita GDP.

To see how this works, suppose that we want to compare two countries, the "US" (which uses dollars and has a million people) and "India" (which uses rupees and has two million people). For simplicity, suppose that these economies only produce computers and education, with the latter only involving the cost of teachers. Assume that the US has 1,000 teachers, each paid $30,000 annually; and sells 10,000 computers annually for $1,000 apiece. As Table 10.1 shows, total GDP will then be $40 million, for a GDP per capita of $40 per year.

Assume further than the exchange rate is 46 rupees per dollar. Computers are tradable goods, and so (roughly) should cost the same everywhere. Thus in India they will cost 46,000 rupees. And let us suppose that teachers in India are paid 36,800 rupees annually (equivalent to $800 annually when converted at the exchange rate). No doubt many Indian teachers would like to move to the US to earn a higher salary, but visa restrictions do not permit this; thus the salary differential persists. As may be seen from Table 10.1, total GDP will be 174.8 million rupees, or 87.4 rupees per capita.

Table 10.1. Computing GDP/capita in purchasing power parity (PPP) terms

	US		India	
	Price ($)	Quantity	Price (Rp)	Quantity
Teachers	30,000	1,000	36,800	1,000
Computers	1,000	10,000	46,000	3,000
GDP	$40 million		Rp. 174.8 million = $4.6 million at 46 Rp/$	
Population	1,000,000		2,000,000	

GDP/capita	$40	Rp.87.4 = $1.90 at 46 Rp/$
GDP/capita, US prices	$40	$16.5
GDP/capita, Indian prices	Rp.496.8	Rp.87.4

Using the exchange rate, Indian GDP is just $1.90 per capita, or 4.75% of the US level. Yet it is clear that this does not do justice to the volume of goods and services produced in India, and so understates India's real GDP relative to that of the US.

One solution would be to value GDP in both countries using US prices. This gives a GDP per capita of $16.50 for India, or 41.25% of the US level of $40. Alternatively one might value GDP in both countries using Indian prices. This gives a GDP per capita of 496.8 rupees for the US; now the Indian level is 17.6% of the US level. In short:

	GDP per capita		
	In US	In India	India/US, %
Using exchange rate, in $	40.00	1.90	4.75
Using US prices ($)	40.00	16.50	41.25
Using Indian prices (rupees)	496.8	87.4	17.6

So although it is clear that using exchange rates to compare GDP per capita is not generally appropriate, we are left with the difficult issue of what common set of prices to use instead. There is no entirely satisfactory answer to this question.

10.3.2 Further methodological considerations

When trying to compare poverty – as opposed to GDP per capita – we have a further problem, which is what basket of goods and services to value. In the above example, we sought to value an average basket of goods and services, but this is not appropriate when are concern is with comparing living standards for the poor.

Conceptually, the answer is not too difficult, as Deaton (2003) emphasizes: to construct a PPP exchange rate that is appropriate for determining a poverty line, take a benchmark consumption basket of the poor in one country and price it directly in the other countries. But the keyword here is "basket:" to construct the exchange rate, a basket will have to be selected and priced in each country, and then aggregated into two indices, namely the price levels for each country. Even single-country price level calculations present some basic dilemmas e.g. the choice of whether to use the base year quantities or current year quantities as weights, how to handle changes in the base year, and classic index number

anomalies. Alatas *et al* (2004) provide a discussion of the various aggregation formulae that can be used to calculate PPP price indices.

These problems are present and magnified for poverty price level measures. Most obviously, consumption bundles of the poor are not identical – even within the same country (a prime and well-studied example is variation in consumption baskets between Indian states; and also within Indonesia). If an economy wide basket is used, then measured weights will be driven by expenditure patterns of non-poor as well as poor. Measured prices may not be applicable for the poor, given substantial differences in product quality.

There is no ideal solution to these problems and the poverty analyst can only work with the best available data while being aware of the limitations. PPPs for global poverty counts first came from the Penn World Tables (PWT), and are now calculated directly by the World Bank as part of ICP benchmarking project. But not all countries have provided ICP data (or at least *recent* ICP data), so their poverty PPPs are imputed using data from countries at similar stages of development. This is usually done by estimating a regression along the lines of

$$Er(PPP)/er = a + b.\ \text{literacy rate} + c.\ \text{food consumption/capita} + \ldots$$

for countries for which PPP computations have been make. The estimated equation is then used to predict the PPP exchange rate for the remaining countries.

Over time, direct PPP measures have been computed for more countries; the added measures have differed substantially from the previously imputed measures, leading to some pronounced changes in the poverty count. So while imputation may be unavoidable, we must recognize that it is a noisy and imperfect procedure.

The underlying basket weights have also been subject to dispute: some argue that we need specific PPPs for the poor, linked more closely to food or nutrition. As it is, PPPs corresponding to different standard macroeconomic aggregates e.g. consumption and GDP, do not always agree with each other.

The earlier Millennium Development Goals poverty counts were based on 1985 PPP exchange rates, and because the price benchmarking data that underlie the 1985 and 1993 counts are completely different, it is probably unwise to look for a simple procedure that would relate the old and the new measures; in particular, the 1993 PPP exchange rates are not simply the 1985 rates updated by the

inflation differential between the US and the relevant country; leaving aside the huge differences in source data, such a procedure would assume no changes in the real exchange rate over this period, an assumption that is not warranted.

10.3.3 Illustrating the importance of exchange rate adjustments

The issues discussed here are usefully illuminated using the World Bank's PovcalNet web interface.[11] For all the household surveys that underlie the global poverty monitoring counts, the Bank has made available estimated distributions for each country for the reference years of the global count. The user can then input any monthly dollar poverty line and any exchange rate to see how the different measures of poverty (headcount, gap, gap squared) are affected (with the $1 per day parameters as the default options). The detailed output also provides the parameters of the Lorenz curve and other relevant statistics.

Exercise: Download the POVCAL.EXE file from the World Bank's PovcalNet site, along with the INDIA.DAT file with data on income distribution for India. Compute the poverty rates (P_0, P_1 and P_2) for India assuming (a) a poverty line of 100 rupees per person per day; (b) a poverty line of 150 rupees per person per day. [Note: You are on the right track if you find that the headcount poverty rate for the poverty line of 100 rupees is 54-55%, depending on the functional form used to approximate the Lorenz curve.]

The case of Uganda is an instructive if extreme example of the measurement volatility induced by the need for PPP exchange rates. Using *PovcalNet* to get the dollar a day headcount for poverty in 1999, we arrive at the figure of 85 percent – high, but not entirely surprising for Sub-Saharan Africa. But note that *PovcalNet* allows the user to see the default parameter settings for this calculation, which includes a PPP exchange rate of about 260 shillings per dollar. This estimate was taken from the PWT version 5.7. For some perspective on this number, consider that the actual nominal exchange rate for 1993 would have been about 1,000 shillings per dollar. In fact, a division by four is not a bad rule of thumb for converting nominal to PPP exchange rates! Thus the PPP adjustment goes in the direction that we expect,

[11] The World Bank provides background information for its global poverty counts at http://www.worldbank.org/research/povmonitor/. Use the "Software" link to get to the POVCAL program, developed by Shaohua Chen, Gaurav Datt and Martin Ravallion for calculating poverty measures from grouped data. The "PovcalNet" link leads to an interactive program that is supposed to allow one to find the World Bank measures of dollar-a-day poverty rates, and permit the user to explore the implications of making other assumptions about the PPP exchange rate.

recognizing that the cost of goods is much lower in Uganda and therefore that the purchasing power of the Ugandan shilling is much higher than the nominal exchange rate would have suggested.

However, the World Bank's own calculations of a PPP exchange rate for Uganda from International Comparison Project data led to a rate of about 84 shillings per dollar. Such a dramatic scaling up of Ugandan purchasing power versus both the nominal exchange and the PWT estimate inevitably scales down the headcount measure of poverty for Uganda against the dollar per day standard, to about 27 percent. At this point of course the discretion of the researcher is a vital component of the exercise; recognizing the unrealism of this finding, the Bank adopted the PWT number and the associated headcount measure instead. In less extreme cases than Uganda, the problem of choosing a PPP exchange rate is arguably more severe, precisely because of the lack of a disqualifying criterion. There are a considerable number of cases where the 1993 PPP estimates of the World Bank and PWT are plausible yet carry with them sizable differences in the estimated number of poor people.

In the absence of any ideal solution, Deaton advocates a common sense approach in which global poverty monitoring stays with the 1993 PPP lines for as long as possible, as the best available data, but checked by local experts and updated with domestic price indices as much as possible. The domestic price indices are less prone to contamination by international price movements of little relevance to the poor. New PPPs would be incorporated only infrequently, and only when there is a clear gain in informational quality. This would reduce volatility in measures due only to statistical revisions, and would provide a consistent basis for the measurement of changes in poverty over time; considerable uncertainty in levels will be a fact of life.

10.4 Survey data and national accounts

Standard measures of poverty are constructed from household budget surveys (HBS). Thus they incorporate the HBS sampling frame and the HBS measure of consumption. The potential problems with this source have been known for a long time but acquired new topicality with research on "pro-poor growth, " which looks at the extent of poverty reduction arising from economic growth. There were some findings for the 1990s that poverty was not falling at the rate we would expect given economic growth – especially in India – yet no evidence of rising measured inequality; so where did the growth "go?"

A typical HBS is designed to be a representative sample (e.g. via the most recent census). But even with an *ex ante* representative sample, the *ex post* sample can be systematically biased. There are actually two problems:

a. The first is non-response bias. There is strong evidence that rich households are less likely to comply with HBS reporting.
b. The second is underreporting bias. It is a general rule that richer households, when they do respond to a survey, are more likely to understate their true income (or expenditure.

These effects probably reflect a desire to conceal income or consumption, as well as the opportunity cost of time.

The first result is that upper income households are underrepresented in the HBS; so headcount *ratio* potentially overstated.

But this is not the end of the story. The macroeconomic national accounts will "see" the transactions of upper income households in their expenditures, so that they are reflected in the key aggregates (national income and/or private consumption). There is a endemic tendency of "average" national accounts consumption to be higher than "average" consumption from the HBS. Furthermore, this gap is not simply an invariant level effect – it is systematically related to the level of development and the growth rate. As a rule of thumb, national accounts consumption exceeds HBS consumption and the divergence increases with economic growth. Deaton (2003) reports that HBS consumption averages 86% of national accounts consumption, and grows about half as rapidly.

One solution to the under-reporting problem is to adjust the HBS figures upwards. Suppose that national consumption, based on extrapolation from HBS data, is $80 million, but national accounts indicate consumption of $90 million. Then one could adjust the HBS data by scaling consumption upwards – i.e. multiplying everyone's reported consumption by 1.125 (=90/80). This would of course reduce the measured poverty rate.

The net effect of under-reporting plus rescaling is unclear, as the following example in Table 10.2 shows. A society has 12 individuals, with incomes as shown here. The true headcount poverty rate is 33.3%. If a rich individual does not respond, the observed poverty rate rises to 36.4%. But if an effort is

made to scale all incomes upwards, then the observed poverty rate would be 18.2%, which represents an overadjustment – essentially because one is adjusting all income categories for underreporting, when the problem is confined to high-income individuals. Underreporting can have a similar effect, in that it would raise the observed poverty rate; however, a rescaling runs the risk of understating poverty, again because the problem of underreporting is not found equally in all income groups, but is a greater problem for high-income individuals.

Table 10.2. Illustrating the effects of response bias and underreporting

Individual #	1	2	3	4	5	6	7	8	9	10	11	12	P0
Response Bias													
True Income/capita	80	100	110	121	190	220	270	310	400	490	620	750	33.3%
Response bias:													
Observed income/cap	80	100	110	121	190	220	270	310	400	490	620		36.4%
Adjust for bias:													
Scaled income/cap	101	126	138	152	239	277	340	390	503	616	780		18.2%
Underreporting													
True Income/capita	80	100	110	121	190	220	270	310	400	490	620	750	33.3%
Underreporting													
Observed income/cap	80	100	108	121	180	220	250	300	400	450	560	600	33.3%
Adjust for bias:													
Scaled income/cap	87	109	117	131	196	239	272	326	435	489	609	652	25.0%

Note: Rescaling for response bias is done by multiplying observed income/capita by the sum of true income per capita (obtained from national accounts data) divided by the sum of observed income per capita. A similar adjustment is made to rescale underreported income.

10.4.1 Why HBS and national accounts total differ

Part of the difference in consumption, as measured by household surveys and by national accounts, reflects different concepts of consumption. The national accounts, but not the HBS, measure the imputed rent on housing, imputed financial intermediation (the saving-lending interest rate gap), and consumption of non-profits. The HBS does provide a better stock-flow linkage than the national accounts (which don't really measure changes in stocks at all), so in principle the HBS allows a separation of "sustainable" consumption from financial income as opposed to asset draw-down, However the practical usefulness of this feature is limited by the fact that households may not report the relevant data correctly. Neither source properly measures consumption of publicly provided goods, access to which plausibly varies with income

Overall though, the consumption of items missed by the HBS is correlated with wealth, both in cross-section and over time. The growth process reinforces the gap, because the national accounts

register the increase in market-based consumption as an economy grows – not all of which represents true economic growth. Similarly, growth shifts activities from the informal to the formal sector. However, and we will return to this point later, it is precisely because the national accounts tend to better reflect the process of wealth accumulation that the HBS measure could still be a "truer" measure of consumption *for the poor.*

10.5 The Debate: Is world poverty falling? If so, how fast?

The current debate about poverty trends is centered on the appropriate measure of consumption. Bhalla (2002) argues that the national accounts measure is more accurate. Thus he measures the distribution of consumption from the HBS, but adjusts the level upwards to correspond to the national accounts. This method results in a much lower level of measured poverty in India, because the divergence between the two sources is severe there, and it predicts much more optimistic trends in poverty reduction than standard projections. In fact, it implies that we already have met, or will soon meet, the MDGs for poverty.

This debate has been conducted in most detail for India, which alone matters a lot for global poverty monitoring because of India's size. However, the issues are germane for other countries as well.

Critics have focused on the fact that Bhalla uses the HBS to get the distribution of consumption but the national accounts to get the mean. He is thus assuming a very specific type of measurement error in the HBS – that it is all in the mean, and none in the variance. Conversely, there is assumption that the national accounts have less measurement error in the mean than the HBS. It is highly unlikely that the measurement errors take exactly this form. In reality, both sources have errors in both mean and variance.

Consider in particular the national accounts estimate of consumption. The fundamental macroeconomic identity gives

GDP	=	C	+	I	+	G	+	X	-	IM
		Consumption		Investment		Gov. spending		Exports		Imports

In most cases, consumption is estimated as a residual, obtained from estimated final production (i.e. GDP) less net exports and government consumption (for which relatively good data exists). For a poor country, the production estimate begins as physical volume (e.g. projected crop yields times estimated crop area), and is then converted to values. This is coupled with a fixed coefficient assumption

for the rest of the production sector (e.g. the allocation of goods to intermediate versus final usage). Notice the multiple entry points for error.

The information on distribution may also be suspect. Practitioners of the national accounts method often use the published income or consumption quintiles (as opposed to the unit-level survey data), which may disguise severe problems in the underlying data. Of course in this respect it is unfair to blame researchers for their use of quintile data, because the unit-level data are often not publicly available.

In any event, the direction of errors induced by the acknowledged flaws in the HBS is not clear-cut. Even if the under-representation of upper incomes causes the HBS to underestimate mean and variance, the "true" distribution may not have a greater mass below the poverty line than the estimated one. However, the National Accounts consumption adjustment, being a pure scale adjustment, clearly reduces the headcount measure.

This seemingly technical debate about data sources has been conducted against the background of an increased results-orientation of international development policy, and specifically an expectation that the cumulative official resource flow to developing countries and the associated initiatives (PRSP/HIPC) should by now have led to an appreciable reduction in poverty. However, countries that have been considered success stories in terms of their pursuit of these new development frameworks (e.g. Uganda) have displayed rates of poverty reduction considered disappointing relative to the per capita GDP growth. Increased research attention has therefore been directed to the issue of "pro-poor growth," namely the extent to which those at the lower end of the income distribution benefit from the growth process. It is the antithesis of this approach to simply ascribe the mean rate of consumption growth to the entire distribution, which assumes that the benefits of growth across the income distribution are neutral.

This assumption is especially problematic for the consumption categories where the two measures don't overlap and where consumption patterns differ systematically across the income distribution. Nevertheless, these arguments cut both ways; if we acknowledge that the HBS data do seem to miss some aspects of growth, it can be questioned whether it is safe to use them to study pro-poor growth. All our information about income distribution comes from the HBS, meaning that we know very little about the distribution of national accounts consumption, yet it is this distribution that may drive public perceptions of the equality of the growth process.

Finally, as with many seemingly tempestuous disputes in economics, some of the apparent differences between different schools of thought weaken in the face of pragmatism. For instance, Karshenas (2004) advocates a hybrid approach to measuring poverty in which survey and national accounts are combined, in recognition of potential errors in both. Similarly, while World Bank researchers have doggedly defended their reliance on survey data, they use national accounts consumption growth data to extrapolate their poverty counts between survey years and to generate forecasts of poverty reduction.

10.6 Conclusion

Let us briefly outline some current directions of research on the linkage between macroeconomic and poverty outcomes. It is now well documented that the elasticity of poverty with respect to growth varies over time and space, prompting investigation of what set of initial conditions makes this elasticity bigger. Furthermore, it is increasingly recognized that while growth remains central to poverty reduction – as the discussion of the Dollar and Kraay results in Chapter 8 made clear – it is not the only way to reduce poverty; particularly in the short- and medium-term inequality also matters.

There are certain policies with modest immediate growth effects but strong poverty impacts, or in the terminology above, policies that lead to a high elasticity of poverty reduction with respect to growth. Prime examples are reforms to security of tenure, micro-credit, and the expansion of basic education. A strict focus on a growth-poverty linkage might overlook such policies.

Poverty analysts are also looking at ways to improving the measurement of consumption, in a way that balances the gain from better source data with the loss of comparability to previous data.

REFERENCES

Alatas, Vivi, Jed Friedman, and Angus Deaton, 2004. "Purchasing power parity exchange rates from household survey data: India and Indonesia," Research Program in Development Studies (Princeton University), Working paper.

Bhalla, Surjit, 2002. *Imagine there is no country: poverty, inequality, and growth in the era of globalization*, Institute for International Economics, Washington DC USA.

Chen, Shaohua and Martin Ravallion, 2001. "How Did the World's Poorest Fare in the 1990s?," *Review of Income and Wealth*, 47(3): 283-300.

Chen, Shaohua and Martin Ravallion, 2004. "How have the world's poorest fared since the early 1980s?" World Bank Development Economics Working Paper No.3341.

Deaton, Angus, 2003. "Measuring Poverty in a Growing World" (or "Measuring Growth in a Poor World"), National Bureau for Economic Research Working Paper 9822 .

Deaton, Angus, 2003. "How to Monitor Poverty for the Millennium Development Goals," *Journal of Human Development*, 2003, 4(3): 353-378.

Karshenas, Massoud, 2004. "Global poverty estimates and the millennium goals: Towards a unified framework" International Labour Organisation Employment Strategy Papers, 2004/5.

Appendix I: Data Introduction

In this course, we will be working extensively with STATA using a subset of information from the Household Survey 1998/99, conducted jointly by Bangladesh Institute of Development Studies (BIDS) and the World Bank. The information was collected at individual, household, and community levels. A description of the data sets and file structure used for these exercises is given below.

File Structure

We will use and generate a lot of files. There are mainly three types of STATA files. Some contain data sets (identified by the suffix `.dta`), others contain STATA programs (identified by the suffix `.do`), and yet others contain a record and output of the work we do in STATA (identified by the suffix `.log`). To keep these files organized, it is useful to create a well-structured set of directories. In what follows we shall assume that you have created the following directory structure:

```
c:\intropov
c:\intropov\data
c:\intropov\dofiles
c:\intropov\logfiles
```

Four files should go into the directory `c:\intropov\data` :

1. `hh.dta` includes 20 household level variables such as household location, household business type, asset ownership, access to service, etc.
2. `ind.dta` includes 10 individual level variables such age, gender, education or schooling, marital status, main activity, working status, occupation, sector, relation to household head, etc.
3. `consume.dta` includes 30 categories of expenditure such as various food and non-food items, rent or rental value of housing, etc.
4. `vprice.dta` includes village level price information on the main food items.

Data Description

	hh.dta
hhcode	household identification number
Thana	thana code (a *thana* is an administrative center comprising a number of villages). It ranges from 1 to 32 as there are 32 thanas.
vill	village code (when combined with thana it uniquely identifies a village). It ranges from 1 to 4 as a maximum four villages are selected from a thana.
region	region code 1 Dhaka (the capital) 2 Chittagong 3 Khulna **4** Rajshahi
weight	Sampling weight for household
distance	distance to nearest paved road (km)
d_bank	distance to nearest commercial/agricultural bank (km)

Toilet	type of latrine used in the household 1 Sanitary 2 non-sanitary
hhelec	if household has electricity 1 yes 2 no
hassetg	household total assets (in *taka*)
famsize	household size
sexhead	gender of household head 1 male 2 female
agehead	age of household head (years)
educhead	years of schooling of household head
hhlandd	land (in *decimals*, i.e. one-hundredth of an acre) owned by household

	Ind.dta	
pid	household member identification number (unique for a household member, so becomes unique in the sample after being combined with household id)	
indsave	individual savings (in *taka*)	
snaghr	non-farm self employment working hours per month	
sagrhr	farm self employment working hours per month	
wnaghr	non-farm wage job working hours per month	
waghr	farm wage job working hours per month	
iemphr	total working hours per month	
rel_hh	code for relation to household head	
	1 Head himself/herself 2 Wife/husband 3 Son/daughter 4 Grandson/granddaughter 5 Father/mother 6 Sister/brother 7Niece/nephew	8 Son-in-law/daughter-in-law 9 Spouse of brother or sister 10 Brother or sister of spouse 11 Father-in-law/mother-in-law 12 Other relatives of head or spouse 13 Servant/maid servant 14 Other _____ (specify)
educ	Years of schooling completed	
sex	Gender	
age	Age (in years)	

	consume.dta
	Ten items have been selected from the survey: rice, wheat, pulses, milk, oil, meat, fish, vegetables, fruits, sugar. Let *X* denote any items, so:
qX	quantity (kg) of item *X* consumed last week
eX	value of item *X* consumed last week (in *taka*)
expfd	household total food consumption per month (in *taka*)
expnfd	household total expenditure on regular non-food items per month (in *taka*)

	vprice.dta
	Eleven price items (vegetables in the above now has two entries: potatoes and other vegetables) were selected from the survey. Again, denote an item by *X*:
pX	village price per kg

Appendix II: STATA Preliminary

STATA is a statistical software package that offers a large number of statistical and econometric estimation procedures. With STATA we can easily manage data and apply standard statistical and econometric methods such as regression analysis and limited dependent variable analysis to cross-sectional or longitudinal data. STATA is widely used by analysts working with household survey data.

1. Getting Started

1.1. Starting STATA

Start a STATA session by double-clicking on the STATA icon in your desktop. The STATA computing environment comprises four main windows. The size and shape of these windows may be moved about on the screen. Their general look and description are shown below:

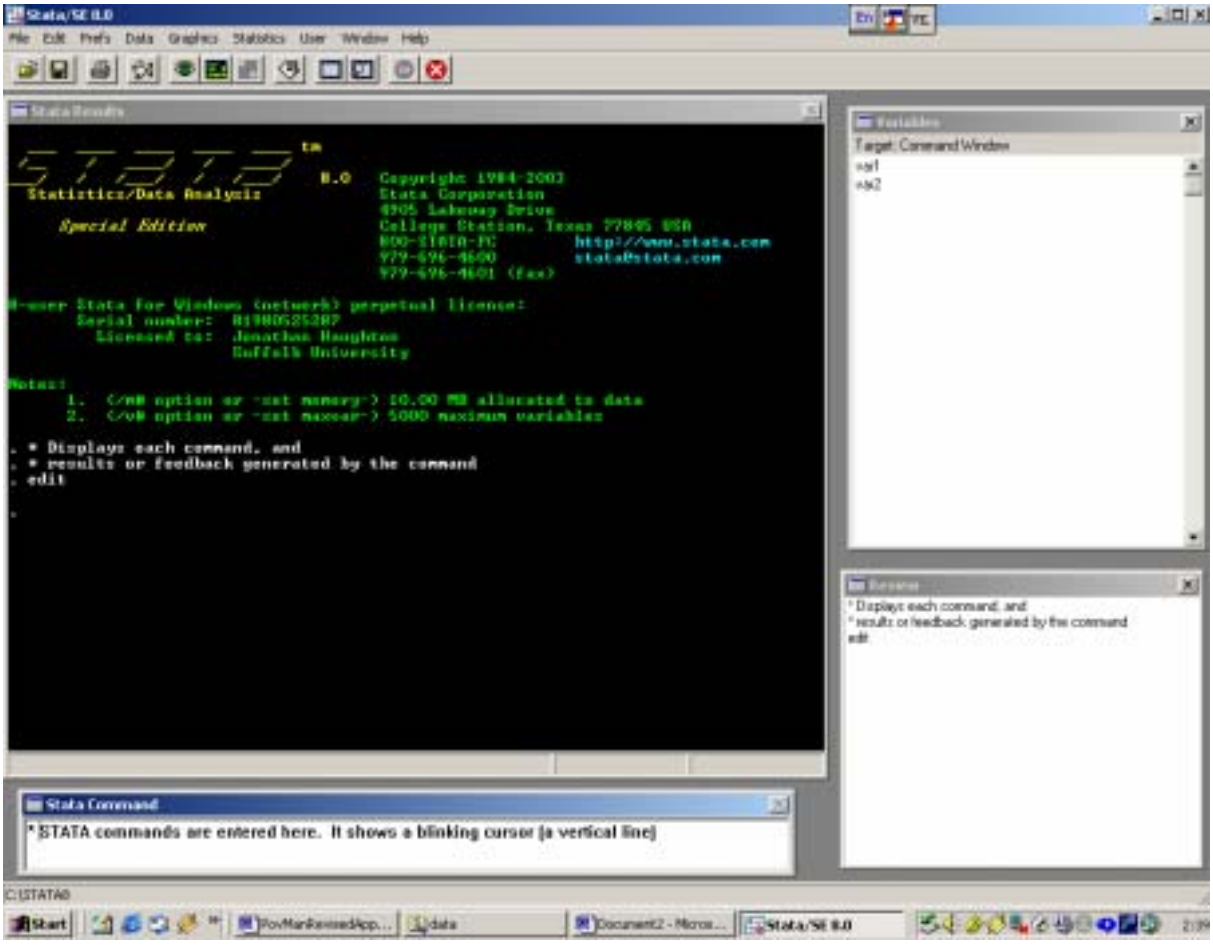

It is useful to have the **`Stata Results`** window be the largest so you can see a lot of information about your commands and output on the screen. In addition to these windows STATA environment has a menu and a toolbar at the top (to perform STATA operations) and a directory status bar at the bottom (that shows the current directory). You can use menu and toolbar to issue different

STATA commands (like opening and saving data files), although most of the time it is more convenient to use the **Stata Command** window to perform those tasks. If you are creating a log file (see below for more details), the contents can also be displayed on the screen; this is sometimes useful if one needs to back up to see earlier results from the current session.

1.2. Opening a Dataset

You open a STATA dataset by entering following command in the **Stata Command** window:

```
use hh
```

(or possibly **use c:\intropov\data\hh,** depending on where STATA has opened up; see below for more details). You can also click on **File** and the **Open** and then browse to find the file you need. STATA responds by displaying the following in the **Stata Results** window:

```
. use hh
.
```

The first line repeats the command you enter and the second line with no error message implies that the command has been executed successfully. From now on, we will show only the **Stata Results** window to demonstrate STATA commands. The following points should be noted:

- STATA assumes the file to be in STATA format with an extension `.dta`. So, typing `hh` is the same as typing `hh.dta`.
- We can only have one data set open at a time in STATA. So if we open another data set `ind.dta`, we will be replacing `hh.dta` with `ind.dta`. Fortunately, STATA does not allow you do to this without warning, unless you type something like **use ind.dta, clear.**
- The above command assumes that the file `hh.dta` is in the current directory (shown by the directory status bar at the bottom). If that is not the case then you can do one of the following two things (assuming current directory is `c:\stata\data` and the file `hh.dta` is in `c:\intropov`):
  - Type the full path of the data file:

```
. use c:\intropov\data\hh
```

  - Or make `c:\intropov` as the current directory and then open the file as before:

```
. cd c:\intropov\data
. use hh
.
```

If the memory allocated to STATA (which is 1,000K or 1M by default) is too little for the data file to be opened, as is typically the case when working with large household survey data sets, we will see an error message like the following:

```
. use hh
no room to add more observations
r(901);

.
```

The third line displays the code associated with the error message. All error messages in STATA have associated codes like this; further explanations are available in the STATA Reference Manuals. In this case we have to allocate more memory to STATA. The following commands allocate 30M to STATA and then try again to open the file:

```
. set memory 30m
[This generates a table with information]
. use hh

.
```

Since the file opens successfully, allocated memory is sufficient. If you continue to get an error message, you can use a larger amount of memory, although this may slow down your computer somewhat. Note that the `set memory` command works only if no data set is open. Otherwise you will get following error message:

```
. use hh

. set memory 10m
no; data in memory would be lost
r(4);
```

You can clear the memory using one of the two commands: `clear` or `drop _all`. The following demonstration shows the first command:

```
. use hh

. set memory 10m
no; data in memory would be lost
r(4);

. clear

. set memory 10m

.
```

1.3 Saving a Data set

If you make changes in an open STATA data file and want to save those changes, you can do that by using the STATA `save` command. For example, the following command saves the `hh.dta` file:

```
. save hh, replace
file hh.dta saved
```

You can optionally omit the filename here (just `save, replace` is good enough). If you do not use the `replace` option STATA does not save the data but issues the following error message:

```
. save hh
file hh.dta already exists
r(602);
```

The `replace` option unambiguously tells STATA to overwrite the pre-existing original version with the new version. If you do NOT want to lose the original version, you have to specify a different filename in the `save` command, for instance using:

```
. save hh1
file hh1.dta saved
```

Notice that there is no `replace` option here. However, if a file named hh1.dta already exists then you have to either use the replace option or use a new filename.

1.4. Exiting STATA

An easy way to exit STATA is to issue the command: `exit`. However, if you have an unsaved data set open, STATA will issue the following error message:

```
. exit
no; data in memory would be lost
r(4);
```

To remedy this problem you can save the data file and then issue the `exit` command. If you really want to exit STATA without saving the data file, you can first clear the memory (using `clear` or `drop _all` command as shown before) and issue the `exit` command. You can also simplify the process by combining two commands:

```
. exit, clear
```

1.5. STATA Help

STATA comes with an excellent multivolume set of manuals. However, the on-computer **help** facility in STATA is extensive and very useful; and if you have access to the Web, an even larger set of macros and other useful information is available.

- From within STATA, if you know which command or keyword you want the help information about, you can the issue the command `help` followed by the command name or keyword. This command only works if you type the full command name or keyword unabbreviated. For example, the following will not work,

```
. help mem
help for mem not found
try help contents or search mem
```

but this will:

```
. help memory

[output omitted]
```

- If you can not recall the full command name or keyword, or you are not sure about which command you want you can use the command `lookup` or `search` followed by the command name or keyword. So following will work:

```
. search mem

[output omitted]
```

This command will list all commands associated with this keyword and display a brief description of each of those commands. Then you can pick what you think is relevant and use "help" to obtain the specific reference.

- The STATA website (**http://www.stata.com**) has excellent help facilities, for example, Online Tutorial, Frequently Asked Questions (FAQ), etc.

1.6. Notes on STATA Commands

Here are some general comments about STATA commands:

- STATA commands are typed in lower case.
- All names, including commands or variable names, can be abbreviated as long as there is no ambiguity. So `describe`, `des` or simply `d` do the same job as there is no confusion.
- In addition to typing, some keystrokes can be used to represent a few STATA commands or sequences. The most important of them are the **`Page-Up`** and **`Page-Down`** keys. To display the previous command in the **`Stata Command`** window, you can press the **`Page-Up`** key. You can keep doing that until the first command of the session appears. Similarly, the **`Page-Down`** key displays the command that follows the currently displayed command in the **`Stata Command`** window.
- Clicking once on a command in the **`Review`** window will put it into the **`Stata Command`** window; double clicking it will tell STATA to execute the command. This can be useful when commands need to be repeated, or edited slightly in the **`Stata Command`** window.

2. Working with data files: looking at the content

From now on, we will mostly list the command(s) and not the results, to save space. To go through this exercise open the `hh.dta` file, as we will use examples extensively from this data file.

2.1. Listing the variables

To see all variables in the data set, use the `describe` command (fully or abbreviated):

```
. describe
```

This command provides information about the data set (name, size, number of observations) and lists all variables (name, storage format, display format, label).

To see just one variable or list of variables use the `describe` command followed by the variable name(s):

```
. d hhcode vill

              storage  display     value
variable name   type   format      label      variable label
------------------------------------------------------------------------------
hhcode          double %7.0f               COMPLETE HOUSEHOLD CODE
vill            float  %2.0f                   VILLAGE CODE
```

As you can see, the `describe` command shows also the variable type, its length and a short description of the variable (if available). The following points should be noted:

- You can abbreviate a list of variables by typing only the first and last variable names, separated by a hyphen (-); the **`Variables`** window shows the order in which the variables are stored. For example, to see all variables from `hhcode` up to `famsize`, you could type:

```
. describe hhcode-famsize
```

- The wildcard symbol (*) is helpful to save some typing. For example, to see all variables that start with `"hh"`, you could type:

```
. describe hh*
```

  You can abbreviate a variable or variable list this way in any STATA command (where it makes sense), not just in `describe`.

2.2. Listing data

To see actual data stored in the variables we use the `list` command (abbreviated as `l`). If you type the command `list` all by itself STATA will display values for all variables and all observations, which is not desirable for any practical purpose (and you may need to use the **Ctrl-Break** combination to stop data from scrolling endlessly across the screen!). We usually want to see the data for certain variables and for certain observations. We achieve this by typing a `list` command with a variable list and with conditions, as shown in the following examples.

The following command lists all variables of the first three observations:

```
. list in 1/3
```

Here STATA displays all observations starting with observation 1 and ending with observation 3. STATA can also display data as a spreadsheet. There are two icons in the toolbar called **`Data Editor`** and **`Data Browser`** (fourth and third from right). By clicking one, a new window will pop up and the data will be displayed as a table, with observations as rows and variables as columns. **`Data Browser`** will only display the data, whereas you can edit data with **`Data Editor`**. The commands `edit` and `browse` will also open the spreadsheet window.

The following command lists household size and head's education for households headed by a female who is younger than 45:

```
. list famsize educhead if (sexhead==0 & agehead<45)
```

The above statement uses two relational operators (`==` and `<`) and one logical operator (`&`). Relational operators impose a condition on one variable, while logical operators combine two or more relational operators. The following list shows the relational and logical operators that are used in STATA:

Relational operators	Logical operators
> (greater than)	~ (not)
< (less than)	| (or)
== (equal)	& (and)
>= (greater than or equal)	
>= (less than or equal)	
!= or ~= (not equal)	

You can use relational and logical operators in any STATA command (where it makes sense), not just in the `list` command.

2.3. Summarizing data

The very useful command `summarize` (which may be abbreviated as `sum`) calculates and displays a few of summary statistics, including means and standard deviations. If no variable is specified, summary statistics are calculated for all variables in the data set. The following command summarizes the household size and education of the household head:

```
. sum famsize educhead
```

Any observation that has a missing value for the variable(s) being summarized is excluded from this calculation by STATA (missing values are discussed later). If we want to know the median and percentiles of a variable, we need to add the `detail` option (abbreviated `d`):

```
. sum famsize educhead, d
```

A great strength of STATA is that it allows for the use of weights. The `weight` option is useful if the sampling probability of one observation is different from another. In most household surveys, the sampling frame is stratified, where the first primary sampling units (often villages) are sampled, and conditional on the selection of primary sampling unit, secondary sampling units (often households) are drawn. Household surveys generally provide weights to correct for the sampling design differences and sometimes data collection problems. The implementation in STATA is straightforward:

```
. sum famsize educhead [aw=weight]
```

Here the variable `weight` has the information on the weight to be given to each observation and `aw` is a STATA option to incorporate the weight into the calculation. We will discuss the use of weights further in the chapter exercises given below.

For variables that are strings, `summarize` will not be able to give any descriptive statistics except that the number of observations is zero. Also, for variables that are categorical (e.g. illiterate = 1, primary education = 2, higher education = 3), it can be difficult to interpret the output of the `summarize` command. In both cases, a full tabulation may be more meaningful, which we will discuss next.

Many times we want to see summary statistics by group of certain variables, not just for the whole data set. Suppose we want to see mean family size, and education of household head, by region. We could use a condition in the sum command (for example, `sum famsize educhead if region == 1 [aw=weight])`, and so on for the other regions, but this is not convenient if the number of categories in the group is large.

There is a simpler solution. First, sort the data by the group variable (in this case, region). You can check this by issuing `describe` command after opening each file. The `describe` command, after listing all

the variables, indicates whether the data set is sorted by any variable(s). If there is no sorting information listed or the data set is sorted by a variable that is different from what you want it to be, you can use the `sort` command and then save the data set in this form. The following commands sort the data set by `region` and show summary statistics of family size and education of household head by region:

```
. sort region
. by region: sum famsize educhead [aw=weight]
```

2.4. Frequency distributions (tabulations)

We often need frequency distributions and cross tabulations. We use the tabulate (abbreviated `tab`) command to do this. The following command gives the regional distribution of the households:

```
. tab region
```

The following command gives the gender distribution of household heads in region 1:

```
. tab sexhead if region==1
```

In passing, note the use of the = = sign here. It indicates that if the regional variable is identically equal to 1, then do the tabulation.

We can use the `tabulate` command to show a two-way distribution. For example, we might want to check whether there is any gender bias in the education of household heads. We use the following command:

```
. tab educhead sexhead
```

To see percentages by row or columns we can add options to the `tabulate` command:

```
. tab region sexhead, col row
```

2.5. Distributions of descriptive statistics (table command)

Another very convenient command is `table`, which combines features of the `sum` and `tab` commands. In addition, it displays the results in a more presentable form. The following `table` command shows the mean of family size, and education of household head, by region:

```
. table region, c(mean famsize mean educhead)
```

```
----------------------------------------
   region |  mean(famsize)  mean(educhead)
----------+-----------------------------
    Dhaka |           5.23               2
Chittagon |           5.82               3
   Khulna |           5.03               3
 Ragfhahi |           5.03               2
----------------------------------------
```

The results are as we expected. But why is the mean of `educhead` displayed as an integer and not a fraction? This is because the `educhead` variable is stored as an integer number and STATA simply truncated numbers after the decimal. Look at the description of this variable.

```
. d educhead
              storage  display     value
variable name   type   format      label      variable label
------------------------------------------------------------------------------
educhead        float  %2.0f                  Education (years) of HH Head
```

We see that `educhead` is a `float` variable: its format (%2.0f) shows that its digits occupy 2 places and it has no digit after the decimal. You can force STATA to reformat the display. Suppose we want it display two places after the decimal, for a three-digit display. The following command shows that and subsequent `table` command:

```
. format educhead %3.2f

. table region, c(mean famsize mean educhead)

----------------------------------------
   region |  mean(famsize)  mean(educhead)
----------+-----------------------------
    Dhaka |           5.23            2.09
Chittagon |           5.82            3.14
   Khulna |           5.03            2.91
 Ragfhahi |           5.03            2.15
----------------------------------------
```

This is much better. Formatting changes only the display of the variable, not the internal representation of the variable in the memory. The `table` command can display up to five statistics, and variables other than the mean (such as the sum or minimum or maximum). It is also possible to display two-way, three-way or even higher dimensional tables.

Here is an example of a two-way table, which breaks down the education of the household head not just by region but also by sex of household head:

```
. table region sexhead, c(mean famsize mean educhead)

----------------------
           |  sex head
           |   1=male
   region |    0     1
----------+-----------
    Dhaka | 3.36  5.37
          | 0.18  2.24
          |
Chittagon | 4.17  6.08
          | 1.50  3.39
          |
   Khulna | 4.18  5.11
          | 1.36  3.05
          |
 Ragfhahi | 3.70  5.13
          | 0.00  2.31
----------------------
```

2.5. Missing Values in STATA

In STATA, a missing value is represented by a dot (`.`). A missing value is considered larger than any number. The `summarize` command ignores the observations with missing values and the `tabulate` command does the same, unless forced to include missing values.

2.6. Counting observations

We use the `count` command to count the number of observations in the data set:

```
. count
  519
.
```

The `count` command can be used with conditions. The following command gives the number of households whose head is older than 50:

```
. count if agehead>50
  159
.
```

3. Working with data files: changing data set

3.1. Generating new variables

In STATA the command `generate` (abbreviated `gen`) creates new variables, while the command `replace` changes the values of an existing variable. The following commands create a new variable called `oldhead`, then set its value to 1 if the household head is older than 32 years and to 0 otherwise:

```
. gen oldhead=1 if agehead>32
(98 missing values generated)

. replace oldhead=0 if agehead<=32
(98 real changes made)
```

What happens here is that, for each observation, the `generate` command checks the condition (whether household head is older than 32) and sets the value of the variable `oldhead` to 1 for that observation if the condition is true, and to missing value otherwise. The `replace` command works in a similar fashion. After the `generate` command, STATA informs us that 98 observations failed to meet the condition and after the `replace` command STATA informs us that those 98 observations have got new values (0 in this case). The following points are worth noting:

- If a `generate` or `replace` command is issued without any conditions, that command applies to all observations in the data file.
- While using the `generate` command, care should be taken to handle missing values properly.
- The right hand side of the = sign in the `generate` or `replace` commands can be any expression involving variable names, not just a value. Thus, for instance, `gen young = (agehead<=32)` would create a variable called `young` that would take on the value of 1 if the head is aged 32 or less (i.e. if the bracketed expression is true) and a value of 0 otherwise.
- The `replace` command can be used to change the values of any existing variable, independently of `generate` command.

STATA provides many useful functions to be used in `generate` and `replace` commands, for example `mean(.)` or `max(.)`. For example, in the `ind.dta` file. the following command calculates the maximum share of employment among four sectors for each household:

```
. gen maxhr=max(snaghr,saghr,wnaghr,waghr)

.
```

An extension of the `generate` command is `egen`. Like the gen command, the `egen` command can create variables to store descriptive statistics like the mean, sum, maximum and minimum or other statistics. For example, an alternative way to create the `maxhr` variable is:

```
. egen maxhr=rmax(snaghr saghr wnaghr waghr)
```

Note the difference in syntax. The more powerful feature of the `egen` command is its ability to create statistics involving multiple observations. For example, the following command creates average individual employment hours for the data set:

```
. egen avgemphr=mean(iemphr)
```

All observations in the data set get the same value for avgemphr. The following command creates the same statistics, this time for males and females separately:

```
. egen avghrmf=mean(iemphr), by(sex)
```

Here observations for males get the value that is average of male employment hours, while observations for females get the equivalent for female employment hours (which is about a tenth as much as for men; check this!).

3.2. Labeling

3.2.1. Labeling variables

You can attach labels to variables to give a description to them. For example, the variable `oldhead` does not have any label now. You can attach a label to this variable by typing:

```
. label variable oldhead "HH Head is over 32"
```

In the `label` command, `variable` can be shortened to `var`. Now to see the new label, type:

```
. des oldhead
```

3.2.2. Labeling Data

We can create other types of labels. To attach a label to the entire data set, which appears at the top of our describe list, try:

```
. label data “Bangladesh HH Survey 1998/99”
```

To see this label, type:

```
. des
```

3.2.3. Labeling Values of variables

Variables that are categorical, like those in `sexhead` (1=male, 0=female), can have labels that help one to remember what the categories are. For example, using **hh.dta**, if we tabulate the variable `sexhead` we see only 0 and 1 values:

```
. tab sexhead
```

To attach labels to the *values* of a variable, we have to do two things. First we have to define a value label. Then we have to assign this label to our variable(s). Using the new categories for `sexhead`:

```
. label define sexlabel 0 "Female" 1 "Male"
. label values sexhead sexlabel
```

Now, to see the labels, type:

```
. tab sexhead
```

If you want to see the actual values of the variable `sexhead`, which is still 0s and 1s, you can add an option to not display the labels we have assigned the values of the variable. For instance, try:

```
. tab region, nolabel
```

3.3. Keeping and Dropping Variables and Observations

We can select variables and observations of a data set by using the `keep` or `drop` commands. Suppose we have a data set with 6 variables: `var1`, `var2`, … , `var6`. We would like to keep a file with only three of them, say `var1`, `var2`, and `var3`. You can use either of the following two commands:

`keep var1 var2 var3` (or `keep var1-var3` if the variables are in this order)
`drop var4 var5 var6` (or `drop var4-var6` if the variables are in this order)

Note the use of a hyphen (`-`) in both commands. It is good practice to use the command that involves fewer variables or less typing (and hence less risk of error!). We can also use relational or logical operators. For example, the following command drops those observations where the head of the household is 80 or older :

```
. drop if agehead>=80
```

And this command keeps those observations where household size is 6 or less:

```
. keep if famsize<=6
```

The above two commands `drop` or `keep` all variables based on the conditions. You cannot include a variable list in a `drop` or `keep` command that also uses conditions. For example, the following command will fail:

```
. keep hhcode famsize if famsize<=6
invalid syntax
r(198);
```

You have use two commands to do the job:

```
. keep if famsize<=6
. keep hhcode famsize
```

You can also use the keyword `in` in a `drop` or `keep` command. For example, to drop the first 20 observations:

```
. drop in 1/20
```

3.4. Producing Graphs

STATA is quite good at producing basic graphs, although considerable experimentation may be needed to produce really beautiful graphs. Version 8 of STATA introduced major changes in the way graphs are programmed. The following command shows the distribution of the age of the household head in a bar graph (histogram):

```
. histogram agehead
```

In many cases, the easiest way to produce graphs is by using the menus; in this case, click on **Graphics** and then on **Histogram** and follow the prompts. An easy way to save a graph is to right click on it, and **Copy** it in order to **Paste** it into Word or Excel.

Here is a command for a scatter plot of two variables. It must typed on a single **`Stata Command`** line.

```
twoway (scatter educhead agehead), ytitle(Education of head) xtitle(Age of head) title(Education by Age)
```

3.5. Combining Data sets

3.5.1. Merging data sets

STATA can only have one data set in memory at a time. However, on many occasions one needs variables that are spread over two or more files and would like to combine those files for the purpose of analysis. For example, we want to see how individual's education varies by the gender of the head of the household. Since the gender variable (`sexhead`) and the individual's education (`educ`) come from two different files (`hh.dta` and `ind.dta`) we have to merge these two files to do the analysis. We want to combine these two files at the household level, so the variable that is used for merging is `hhcode` (this is the *merge variable*). Before merging is done, both files must be sorted by the merge variable. The following command opens, sorts and saves the `ind.dta` file:

```
. use ind,clear
. sort hhcode
. save, replace
```

Once both data sets have been sorted, they can be merged, as follows:

```
. use hh, clear
. sort hhcode
. merge hhcode using ind
```

In the context `hh.dta` is called the *master file* (this is the one that remains in the memory before merging) and `ind.dta` is called the *using file*. To see how the merge operation went, we type the following command:

```
. tab _merge

     _merge |      Freq.     Percent        Cum.
------------+-----------------------------------
          3 |       2767      100.00      100.00
------------+-----------------------------------
      Total |       2767      100.00
```

The variable `_merge` is created by STATA after each merge operation and it can have three possible values:

1 shows those observations from the master file that could not be merged
2 shows those observations from the using file that could not be merged
3 shows those observations that were successfully merged.

The total number of observations in the resulting data set is the sum of these three `_merge` frequencies. A possible candidate for `_merge`=1 would be an observation in the `hh.dta` file with a `hhcode` value that cannot be found in the `ind.dta` file. Similarly if the `ind.dta` file has a `hhcode` that is not found in the `hh.dta` file, that observation will appear with `_merge`=2. In the above example, however, each household in `hh.dta` file has an exact match in the `ind.dta` file and that is why we got `_merge`=3 and not 1's or 2's. If you keep only the matched observations, you can do that by this command: `keep if _merge==3`. After checking the `_merge` variable, it is good practice to drop it from the dataset, using `drop _merge` or simply `drop _m`. Once we have merged the data sets we can go ahead with our analysis:

```
. sort sexhead
. by sexhead: sum educ

-> sexhead = 0

    Variable |     Obs        Mean   Std. Dev.     Min      Max
-------------+--------------------------------------------------
        educ |     179    2.329609   3.424591        0       14
-------------------------------------------------------------------------------
-> sexhead = 1

Variable |     Obs        Mean   Std. Dev.       Min        Max
-------------+--------------------------------------------------
        educ |    2588    2.340417   3.324081        0         16
```

The result shows that there is not much difference in education by the gender of household head.

Notice that in order to show the results by `sexhead` variable, we first have to sort the data by that variable, otherwise we will get an error message.

3.5.2. Appending data sets

Consider what would happen in the above merging scenario if we have `_merge`=1 and 2 only but no 3s. This can happen if the individual data in **ind.dta** come from households that are completely different from the households in **hh.dta**. In this case resulting number of observations after merging is the sum of observations in the two files (observations with `_merge=1` + observations with `_merge=2`). STATA in this case would actually append the two data sets; however variables that are only included in one file will

have missing values for all observations from the other file.. Although this is not what we intend in the above example, appending is necessary when we need to combine two data sets that have same (or almost same) variables but are mutually exclusive. For example, suppose we had four regional versions of the `hh.dta` file: `hhdhak.dta` (has households only from Dhaka region), `hhkhul.dta` (has households only from Khulna region), `hhraj.dta` (has households only from Rajshahi region) and `hhchit.dta` (has households only from Chittagong region). These data files have same variables as does `hh.dta` but represent four distinct sets of households. To combine them to get a overall data set of the whole country we use the `append` command:

```
. use hhdhak
. append using hhkhul
. append using hhraj
. append using hhchit
```

At this stage we have a data set in the memory that has household information from all four regions. If we need this data set for subsequent use we should save it after arranging it in a defined order (say, sorting by `hhcode`).

4. Working with .log and .do files

STATA can work interactively, which is very helpful in debugging commands and in getting a good "feel" for the data. You type one command line each time and STATA processes that command, displays the result (if any), and waits for the next command.

However, often one wants to save the results, and perhaps print them out. This is done by creating a `.log` file. Such a file is created by issuing a `log using` command and closed by a `log close` command; all commands issued in between, as well as corresponding output (except graphs) are saved in the `.log` file. Let us go back to the example in Section 3.5.1. Assume that we want to save only the education summary by household gender, not the merging outcomes. Here are commands:

```
. log using educm.log

. by sexhead:sum educ

-------------------------------------------------------------------------------
-> sexhead = 0

    Variable |     Obs        Mean   Std. Dev.       Min      Max
-------------+--------------------------------------------------------
        educ |     179    2.329609   3.424591         0       14
-------------------------------------------------------------------------------
-> sexhead = 1

Variable |     Obs        Mean   Std. Dev.       Min        Max
-------------+--------------------------------------------------------
        educ |    2588    2.340417   3.324081      0         16
.log close
```

What happens here is that STATA creates a text file named `educm.log` in the current directory and saves the summary output in that file. If you want the .log file to be saved in a directory other than current directory, you can specify full path of the directory in the `.log` creation command. You can also use the **File** button, followed by **Log** and **Begin**.

If a `.log` file already exists, you could either replace it, with `log using educm.log, replace`, or append new output to it with `log using educm.log, append`. If you really want to keep the existing `.log` file unchanged, then you can rename either this file or the file in the `.log` creation command. In a `.log` file if you want to suppress portion of it, you can issue `log off` before that portion and `log on` command after that. You have to close a `.log` file before opening a new one, otherwise you will get an error message.

If you find yourself using same set of commands repeatedly, you can save those commands in a file and run them together whenever you need it. These command files are called `.do` files, and are the STATA equivalent of "macros". There are at least three good ways to create `.do` files:

1. Simply type the commands into a text file; label it `educm.do` (the `.do` suffix is important); and run the file using `do educm` in the **`Stata Command`** window.
2. Right click anywhere in the **`Review`** window; this will save all the commands that were used interactively. The file in which they were saved can be edited, labeled, and used as a `.do` file.
3. Use STATA's built-in `.do` editor. It is invoked by clicking on the icon (the fifth from the right, at the top of the page). Commands may then be typed into the editor. These commands may be run by highlighting them and running them using the appropriate icon (the second from the right) within the `.do` editor. With practice this becomes a very quick and convenient way to work with STATA.

Here is an example of a `.do` file:

```
log using educm.log
use ind
sort hhcode
save, replace
use hh
merge hhcode using ind
tab _merge
sort sexhead
by sexhead:sum educ
log close
```

The main advantages of using `.do` files instead of typing commands line by line are replicability and repeatability. With a `.do` file, one can replicate results that were worked on weeks or months before. And `.do` files are especially useful when sets of commands need to be repeated – for instance, with different data sets, or groups.

There are certain commands that are useful in a .do file. We will discuss them from the following sample `.do` file:

```
*This is a STATA comment that is not executed
/*****This is a do file that shows some very useful
 commands used in do files. In addition, it creates a
log file and uses some basic STATA commands    ***/

#delimit ;
set more 1;
drop _all;
cap log close;
log using c:\intropov\logfiles\try1.log, replace;

use c:\intropov\data\hh.dta ;
describe ;
list in 1/3 ;
list hhcode famsize educhead if sexhead==2 & agehead<45;
summarize famsize;
summarize famsize, detail;
sum famsize educhead [aw=weight], d;
tabulate sexhead;
tabulate educhead sexhead, col row chi;
tabulate educhead, summarize(agehead);
label define sexlabel 1 “MALE” 2 “FEMALE”;
label values sexhead sexlabel;
tabulate sexhead;
label variable sexhead "Head Gender";
use c:\intropov\data\hh.dta;
sort hhcode;
save temp, replace;
use c:\intropov\data\consume.dta, clear;
sort hhcode ;
#delimit cr
merge hhcode using temp
tabulate _merge
keep if _merge==3
drop _merge
log close
```

The first line in the file is a comment. STATA treats any line that starts with an asterisk (*) as a comment and ignores it. You can write multi-line comment by using a forward slash and asterisk (/*) as the start of the comment and end the comment with an asterisk and forward slash (*/). Comments are very useful for documentation purpose and you should include at least following information in the comment of a do file: general purpose of the do file, and last modification time and date. You can include comments anywhere in the do file, not just at the beginning.

#delimit ; By default STATA assumes that each command is ended by the carriage return (ENTER key press). If, however, a command is too long to fit in one line you can spread it over more than one line. You do that by letting STATA know what the command delimiter would be. The command in the example says that a command is ended by a semicolon (;). Every command following the delimit command has to end with a ; until the file ends or another #delimit cr command appears which makes carriage return again the command delimiter. Although for this particular .do file we don't need to use the #delimit command it is done to explain the command.

`set more 1` STATA usually displays results one screenful at a time, and waits for the user to press any key. But this would soon become a nuisance if, after setting a .do file to run, you have to press a key every time this happens until the program ends. This command displays the whole output, skipping page after page automatically.

`drop _all` This command clears the memory.

`cap log close` This command closes any open .log file. If no log file is open, STATA just ignores this command.

Exercise: Open the do-file editor and type the above code into it, then save it as `c:\intropov\dofiles\try.do`. Click the "Do current file" icon and switch to the "STATA Results" window. When you see "end of do file", open `c:\intropov\logfiles\try11.log` in Word and check the results.

Follow-up practice

Now let's do some practice using all three datasets. Remember, do not overwrite these three data files.

1. Generate a new variable, `agegroup`, which categorizes individuals according to their age. For example, assign 1 to `agegroup` if the person is less older than 30. You can make your own categories as you consider appropriate. Label this variable and its categorical values, and tabulate it.
2. Calculate the sex ratio of the sampled population, and the labor participation rates for both men and women.
3. Count the number of children younger than 15 and the number of elders older than 65. Compare the mean per capita staple food consumption (in kg) for households with no children, one child, and two or more children.
4. Calculate the mean per capita food consumption for those households whose heads are aged between 30 and 39. Compare it to the mean per capita food consumption for those whose heads are aged between 50 and 59.
5. Report the mean and median per capita food consumption for each educational level of household head.
6. Calculate the food share in total household expenditure and compare the mean food share for households headed by men with that of households headed by women.
7. Tabulate mean household size and mean educational level by region and area.

Appendix 3: Exercises

Introduction

Working with household data sets requires a solid mastery of appropriate statistical and data management software, such as Stata or SPSS. This mastery comes from learning by doing. We have found that students who work though the exercises in this appendix acquire the necessary mastery, and are ready to tackle almost any challenge in working with household data. The exercises build on one another, so they should be done in the order given, and each completed fully before proceeding to the next one.

Before beginning these exercises, it is important to prepare the data as set out in Appendix 2. If you are new to Stata, you will want to work though Appendix 2; if you once knew Stata, and have forgotten the details, a quick skim of Appendix 2 should suffice to bring back the fond memories.

Exercise 1 (Chapter Two)

We first need to construct the data set that we will use in the later exercises.

1.1 Household Characteristics

Open `c:\intropov\data\hh.dta`, which consists of household level variables. Answer the following questions:

i. How many variables are there? _____

ii. How many observations (households) are there? _____

iii. There are four regions. Household characteristics may vary by regions. Fill in the following table (Hint: use the `table` command).

	Dhaka	Chittagong	Khulna	Rajshahi
Total number of households	_____	_____	_____	_____
Total number of population	_____	_____	_____	_____
Average distance to paved road	_____	_____	_____	_____
Average distance to nearest bank	_____	_____	_____	_____
% Household has electricity	_____	_____	_____	_____
% Household has sanitary toilet	_____	_____	_____	_____
Average household assets	_____	_____	_____	_____
Average household land holding	_____	_____	_____	_____
Average household size	_____	_____	_____	_____

iv. Are the sampled households very different across regions?

v. The gender of the head of household may also be associated with different household characteristics:

	Male-headed households	Female-headed households	
Average household size	_____	_____	
Average years of schooling of head	_____	_____	
Average age (years) of head	_____	_____	
Average household assets (taka)	_____	_____	
Average household land holding (acres)	_____	_____	[Careful!]

[*For consideration:* How many decimal places should one report? As a general rule, do not provide spurious precision. Reporting the average household size as 5.35368 gives a false impression of accuracy; but reporting the size as 5 is too blunt. In such cases, 5.4 or 5.35 would be more appropriate, and is accurate enough for almost all uses.]

vi. Are the sampled households headed by males very different from those headed by females?

1.2 Individual Characteristics

Now open `c:\intropov\data\ind.dta`. This file consists of information on household members. Merge this data with the household level data (`hh.dta`) – see Appendix 2 if you need a refresher on merging – and answer the following questions for individuals *who are 15 years old or older*:

i. Regional variation	Dhaka	Chittagong	Khulna	Rajshahi
Average years of schooling	_____	_____	_____	_____
Gender ratio (% of household that is female)	_____	_____	_____	_____
% Working population (with positive working hours)	_____	_____	_____	_____
% Working population working on a farm	_____	_____	_____	_____

ii. Are the sampled individuals very different across regions?

iii. We now examine some gender differences:

	For males	For females
Average schooling years(age>=5)	_____	_____
Average schooling years(age<15)	_____	_____
Average age	_____	_____
% Working population (with positive working hours)	_____	_____
% Working population working on a farm	_____	_____
Average working hours per month	_____	_____
Average working hours on farm, per month	_____	_____
Average working hours off farm, per month	_____	_____

iv. Are the characteristics of the sampled women very different from those of the sampled men?

1.3 Expenditure

Open `c:\intropov\data\consume.dta.` It has household level consumption expenditure information. Merge it with `hh.dta`.

i. Create three variables: per capita food expenditure (let's call it `pcfood`), per capita nonfood expenditure (call it `pcnfood`) and per capita total expenditure (call it `pcexp`). Now let's look at the consumption patterns.

Average per capita expenditure:

	pcfood	**pcexp**
by region:		
Whole	_____	_____
Dhaka region	_____	_____
Chittagong region	_____	_____
Khulna region	_____	_____
Rajshahi region	_____	_____
by gender of head:		
male-headed households	_____	_____
female-headed households	_____	_____
by education level of head:		
head has some education	_____	_____
head has no education	_____	_____
by household size:		
Large household (>5)	_____	_____
Small household (<=5)	_____	_____
by land ownership:		
Large land ownership (>0.5/person)	_____	_____
Small land ownership or landless	_____	_____

Summarize your findings on per capita expenditure comparison

ii. Now add another measure of household size, which takes into account the fact that children consume less than adults. Assume that children (aged <15) will be weighted as 0.75 of an adult. For instance, a household consisting of a couple with one child aged at 7 is worth 2.75 on this adult-equivalence scale, instead of 3. Go back to the `ind.dta` and create this variable (let's call it `famsize2`), then merge the revised file with the household data and the consumption data files. Create per adult equivalent expenditure variables (let's call them `pafood` and `paexp`) and repeat the above exercise.

Average per capita expenditure:

	pcfood	**pcexp**

by region:

Whole	_____	_____
Dhaka region	_____	_____
Chittagong region	_____	_____
Khulna region	_____	_____
Rajshahi region	_____	_____
by gender of head:		
male-headed households	_____	_____
female-headed households	_____	_____
by education level of head:		
head has some education	_____	_____
head has no education	_____	_____
by household size:		
Large household (>5)	_____	_____
Small household (<=5)	_____	_____
by land ownership:		
Large land ownership (>0.5/person)	_____	_____
Small land ownership or landless	_____	_____

Compare your new results with those of per capita expenditure. In analyzing poverty, is it better to use adult-equivalents?

iii. Besides looking at the mean or the median value of consumption, we can also easily look at the whole distribution of consumption using `scatter`. The following plots the cumulative distribution function curve of per capita total expenditure.

```
. cumul pcexp, gen(pcexpcdf)

. twoway scatter pcexpcdf pcexp if pcexp<20000, ytitle("Cumulative
Distribution of pcexp") xtitle("Per Capita total expenditure") title("CDF
of Per Capita Total Expenditure") subtitle("Exercise 1.3") saving(cdf1,
replace)
```

The `cumul` command creates a variable called `pcexpcdf` that is defined as the empirical cumulative distribution function (cdf) of `pcexp`; in effect it sorts the data by `pcexp`, and

creates a new variable that cumulates and normalizes `pcexp`, so that its maximum value is 1. To explore the variable, try

```
list pcexp pcexpcdf in 1/10
sort pcexp
list pcexp pcexpcdf in 1/10
list pcexp pcexpcdf in –10/-1
```

Then use the code shown here to graph the cdf. Feel free to experiment with the `scatter` command. The graph is also saved in a file called `cdf1.gph`. When you want to look the graph later, you just need to type "`graph use cdf1`".

The cumulative distribution function curve of a welfare indicator can reveal quite a lot of information about the poverty and inequality. For example, if we know the value of a poverty line, we can easily find the corresponding percentage value of people below the line. Suppose the poverty line is 5,000. Then the command

```
sum pcexpcdf if pcexp<5000
```

will give the poverty rate (under the "max" heading).
[For consideration: Why is the mean not the appropriate measure of poverty here?]

iv. Keep `pcfood pcexp pafood paexp famsize2 hhcode`, merge with `hh.dta`, sort by `hhcode`, and save as `pce.dta` in the `c:\intropov\data` directory.

1.4 Household Weights

In most household surveys, observations are selected through a random process, but different observations may have different probabilities of selection. Therefore, we need use weights that are equal to the inverse of the probability of being sampled. A weight of w_j for the *j*th observation means, roughly speaking, that the *j*th observation represents w_j elements in the population from which the sample was drawn. Omitting sampling weights in the analysis usually gives biased estimates, which may be far from the true values (see chapter 2).

Various post-sampling adjustments to the weights are sometimes necessary. A household sampling weight is provided in the `hh.dta`. This is the right weight to use when summarizing data that relates to households.

However, often we are interested in the individual, rather than the household, as the unit of analysis. Consider a village with 60 households; thirty households have five individuals each (with income per capita of 2,100), while the other thirty households have 10 individuals each (with income per capita of 1,200). The total population of the village is 450. Now suppose we take a 10% random sample of households, picking three five-person households and three 10-person households. We would estimate the mean income per capita to be 1,650. While this properly reflects the nature of *households* in the village, it does not give information that is representative of *individuals*: the village has 150 people in 5-person households and 300 people in 10-person households. Weighted by individuals, per capita income in this village is in fact 1,500 [Try the calculation!]. Such computations can be done easily in STATA.

In estimating individual-level parameters such as per capita expenditure, we need to transform the *household* sample weights into *individual* sample weights, using the following STATA commands:

```
. gen weighti = weight*famsize
```

```
. table region [pweight=weighti], c(mean pcexp)
```

STATA has four types of weights: `fweight`, `pweight`, `aweight`, and `iweight`. Of these, the most important are:

- ***Frequency weights*** (`fweight`), which indicate how many observations in the population are represented by each observation in the sample, must take integer values.
- ***Analytic weights*** (`aweight`) are especially useful when working with data that contain averages (e.g. average income per capita in a household). The weighting variable is proportional to the number of persons over which the average was computed (e.g. number of members of a household). Technically, analytic weights are in inverse proportion to the variance of an observation (i.e. a higher weight means that the observation was based on more information and so is more reliable in the sense of having less variance).

Further information on weights may be obtained by typing `help weight.`

Now let's repeat some previous estimation with the newly-created weights:

	Dhaka	Chittagong	Khulna	Rajshahi
Average household size	_____	_____	_____	_____
Average per capita food expenditure:	_____	_____	_____	_____
Average per capita total expenditure:	_____	_____	_____	_____

Are the weighted averages very different from unweighted ones?

1.5 The effects of clustering and stratification

If the survey under consideration has a complex sampling design, then the standard errors of estimates (and sometimes even the means) will be biased if one ignores clustering and stratification.

Consider the following typical case of a multistage stratified random sample with clustering.

i. First one divides the country into regions (the ***strata***), and picks a sample size for each region. Note that it is perfectly legitimate to sample some regions more heavily than others; indeed one would typically want to sample a sparsely populated heterogeneous region more heavily (e.g. one person per 300) than a densely populated, homogeneous region (e.g. one person per 1,000).

ii. Within each region one randomly picks communes, where the probability that a commune is picked depends on the population of the commune; in this case the commune is the primary sampling unit (the ***psu***). Within the commune one may survey households in a cluster – for instance picking 20 households in a single village. Cluster sampling is widespread, because it is much cheaper than taking a simple random sample of the population. Let us assume that someone has also computed a weight variable (***wt***) that represents the number of households that each representative household "represents;" thus the weight will be small for over-sampled areas, and larger for under-sampled areas.

STATA has a very useful set of commands designed to deal with data that have been collected from multistage and cluster samples surveys. First one needs to provide information on the structure of the survey using the `svyset` commands. Using our example we would have

```
svyset [pweight=weighti], strata(region) psu(thana) clear(all)
```

where `region` is a variable that indicates the regions.[12] Having set out the structure of the survey, one may now use `svymean` to give estimates of population means and their correct standard errors; and `svyreg` to perform linear regression, taking into account of survey design. Other commands include `svytest` (to test whether a set of coefficients are statistically significantly different from zero) and `svylc` (to test linear combinations, such as the differences between the means of two variables). Repeat the above exercise (1.4) and compare the results.

	Dhaka	Chittagong	Khulna	Rajshahi
Average household size	_____	_____	_____	_____
Average per capita food expenditure:	_____	_____	_____	_____
Average per capita total expenditure:	_____	_____	_____	_____

Are the new weighted averages, adjusted for clustering and stratification, very different from the unweighted ones?

[12] These commands were substantially revised in STATA version 8, and the syntax differs significantly from earlier versions of STATA.

Exercise 2 (Chapter Three)

In order to compare poverty measures over time, it is important that the poverty line itself represent similar levels of well-being over time and across groups. Three methods have been used to derive poverty lines for Bangladesh: direct caloric intake, food energy intake and cost of basic needs.

The following table gives a nutritional basket – in per capita terms – considered minimal for the healthy survival of a typical adult in a family in rural Bangladesh.

Food items	Per capita normative daily requirements		Average rural consumer price (taka/kg)
	Calorie	Quantity (gm)	
Rice	1,386	397	15.19
Wheat	139	40	12.81
Pulses	153	40	30.84
Milk (cow)	39	58	15.90
Oil (mustard)	180	20	58.24
Meat (beef)	14	12	66.39
Fish	51	48	46.02
Potatoes	26	27	8.18
Other vegetables	26	150	38.30
Sugar	82	20	30.49
Fruit	6	20	28.86
Total	**2102**	**832**	

2.1 Direct Caloric Intake

The direct caloric intake method considers any household not meeting the nutritional requirement of 2,102 Calories per day per person as poor.[13] For this method, we need to know the quantity of every food item consumed by households and its calorie content. Then we calculate the total calorie content of the food actually consumed and derive an equivalent daily caloric intake per capita for each household. `c:\intropov\data\consume.dta` includes the quantity of 10 food items consumed ("potatoes" and "other vegetables" listed above are combined into one item called "vegetables" in the survey; assume that the total per capita daily calorie requirement of this combined item is 52 and the quantity is 177 gm).

1. Use the quantity information from the data set and the calorie content information from the above table to calculate each household's per capita caloric intake (in Calories per day). [Hint: The unit in the data set is kg per week, and this needs to be converted into gm per day.]
2. Create a new variable `cpcap` to store this caloric intake variable. Now identify the households for which `cpcap` is less than 2,102. These households are considered "poor" based on the direct caloric intake method. Create a variable `directp` that equals 1 if the household is poor and 0 otherwise. What percentage of people are poor by this method?

	Bangladesh	Dhaka	Other regions
% poor using direct caloric intake method	**58.8**		

[13] A calorie is the energy required to heat one gram of water by one degree Celsius. A Calorie is 1,000 calories.

2.2 Food-Energy Intake

The food-energy intake method finds the value of per capita total consumption at which a household can be expected to fulfill its caloric requirement, and determines poverty based on that expenditure. Note that this expenditure automatically includes an allowance for both food and non-food, thus avoiding the tricky problem of determining the basic needs for those goods. It does not need price data either. But as explained in Chapter 3, this method can also give very misleading results.

A simple way to implement this method is to rank households by their per capita caloric intakes and calculate the mean expenditure for the group of households who consume approximately the stipulated per capita caloric intake requirement. Proceed as follows:

i. Merge `cpcap` with `hh.dta` and calculate the average `pcexp` for the households whose per capita calorie intake is within 10% minus/plus range of 2,102 (see code in box below).

ii. Call the average value `feipline` and identify the households for whom `pcexp` is less than `feipline`. These households are considered "poor" based on the food-energy intake method. Create a variable `feip` that equals 1 if the household is poor and 0 otherwise.

```
. sum pcexp [aw=weighti] if cpcap<2102*1.1 & cpcap>2102*.9
. gen feipline = r(mean)
. gen feip = (pcexp <= feipline)
```

Technical aside: Note that STATA commands that report results also save the results so that other commands can subsequently use those results; "r-class" commands such as `summarize` save results in `r()` in version 6.0 or higher. After any r-class commands, if you type "`return list`", STATA will list what was saved. [Try it!]

Another group – "e-class" commands such as `regress` – save results in `e()` and "`estimates list`" will list saved results. For example, `e(b)` and `e(V)` store the estimates of coefficients and the variance-covariance matrix, respectively. To access coefficients and standard errors, there is an easier way. `_b(varname)` or `_coef(varname)` contains the coefficient on `varname` and `_se(varname)` refers to the standard error on the coefficient.

iii. What percentage of people are poor by this method?

	Bangladesh	Dhaka	Other regions
% poor using food energy intake method	_____	_____	**67.9**

iv. *Challenge:* a more sophisticated method is to regress per capita total expenditure on per capita calorie intake and then predict the expected per capita expenditure at 2,102 kcal level. Try this!

```
. regress pcexp cpcap [aw=weighti]
. gen feipline=_b[_cons] + _b[cpcap]*2102
```

v. Should there be separate regression for each region?

2.3 Cost of Basic Needs

The idea behind the cost of basic needs method is to find the value of consumption necessary to meet minimum subsistence needs. Usually it involves a basket of food items based on nutritional requirements and consumption patterns, and a reasonable allowance for non-food consumption.

i. According to the above basket and the average rural consumer prices, how much money does a household of four need each day to meet its caloric requirements?

ii. One way to derive the non-food allowance is simply to assume a certain percentage of the value of minimum food consumption. How much annual total expenditure does a family of four need if it is to avoid being poor, assuming that non-food expenses amount to 30 percent of food expenses?

iii. `vprice.dta` gives village-level price information on all 11 food items. Therefore, we can actually calculate a food poverty line (call it `foodline`) and a total poverty line (call it `cbnpline`) for each village using the cost of basic needs method and merge this variable with `pce.dta.` [Hint: Here we need to `sort` both data sets and `merge` by `thana vill`]. Do this, and create a variable `cbnp` that equals 1 for the poor and 0 for the non-poor.

iv. What percentage of people are poor by this method?

	Bangladesh	Dhaka	Other regions
% poor by cost of basic needs method	_____	_____	_____

v. The percentage of people in poverty varies according to the three methods. Which method do you consider to be most suitable here? Why?

---

vi. Keep all imputed poverty lines and poverty indicators, merge with `pce.dta`, and save the file as `final.dta`.

Exercise 3 (Chapter Four):

3.1 A Simple Example

In STATA, open the data file `example.dta` and browse the data using STATA "Data Browser". You should see a spreadsheet listing information exactly as presented in the following table.

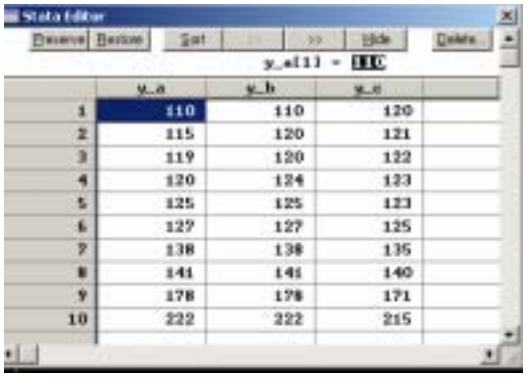

	y_a	y_b	y_c
1	110	110	120
2	115	120	121
3	119	120	122
4	120	124	123
5	125	125	123
6	127	127	125
7	138	138	135
8	141	141	140
9	178	178	171
10	222	222	215

The data consists of information on consumption by all the individuals in three countries (A, B and C). Each country has just 10 residents.

1. Summarize the consumption level for each of the three countries:

2. Assuming a poverty line of 125, calculate the following poverty rates for each country:

Country	A	B	C
a. Using the headcount index:	_____	_____	_____
b. Using the poverty gap index:	_____	_____	_____
c. Using the squared poverty gap index:	_____	_____	_____

   [Hint: The relevant formulas are provided in Chapter 4. Try programming the results in STATA, rather than doing the computations by hand or using Excel.]

3. Which country has the highest incidence of poverty? Justify your answer.

3.2 Poverty Measures for Rural Bangladesh 1999

Now let's work with the per capita food expenditure and the per capita total expenditure (`pcfood` and `pcexp` in `c:\intropov\data\final.dta`) that we have created in Exercise 1, and use `cbnpline` (the cost of basic needs poverty line that we derived in Exercise 2).

Technical note: Although it is possible to program the calculation of different measures of poverty, it is simpler use programs that have been written by others. In STATA these programs are

known as `.ado` programs. The basic version of STATA comes with a large library of such programs, but for specialized work (such as computing poverty rates) it is usually necessary to install `.ado` programs that have been provided on a diskette or obtained on the Web.

For computing poverty rates, and their accompanying standard errors, we have provided `FGT.ado` , which is based on `poverty.ado` written by Philippe Van Kerm; the standard error calculation follows Deaton (1997). The `FGT.ado` file should be put in your working directory; or into a directory given by c:\ado\plus\f (which you may need to create for this purpose). We have also provided two other useful `.ado` programs, `SST.ado` (for computing the Sen-Shorrocks-Thon measure poverty) and `Sen.ado` (for computing the Sen index of poverty). Other `.ado` programs are available on the Internet; for an example, and how to access them, see section 3.3 below.

`FGT.ado` can calculate the head count index (or FGT(0)), the poverty gap index (or FGT(1)), and the squared poverty gap index (or FGT(2)). For example,

```
. FGT y, line(1000) fgt0 fgt1 fgt2
```

will calculate the headcount ratio, the poverty gap ratio, and squared poverty gap index using a poverty line of 1000 and welfare indicator `y`. Be careful: the command is case sensitive, and in this case `FGT` must be written in capital letters. After `line`, the brackets must contain a number. Instead of typing all three measures one could specify *`all`* option, or just some of the measures. If one also types `sd`, then the command will also give standard errors for the estimates, which is very useful in determining the size of sampling error.

The above command works when there is a single poverty line. However, some researchers prefer to compute different poverty lines for each household (as a function of the household size, local price level, etc.). Assume that these tailor-made poverty lines are in a variable called `povlines`. Now the appropriate command becomes

```
. FGT y, vline(povlines) fgt0 fgt1 fgt2 sd
```

You can specify conditions, range and weights with these commands. For example, the following command calculates the headcount ratio for Dhaka region based on a poverty line of 3000.

```
. FGT pcexp [aw=weighti] if region==1, line(3000) fgt0
```

`Sen.ado` and SST.ado calculate the Sen index and SST index respectively. The syntax follows the same format, but does not compute standard errors. So, for example, one could use:

```
. Sen y, line(1000)
. SST y, line(1000)
```

Now we are ready to turn to the measurement of poverty using the data from the BHES 1991/2.

1. Compute the five main measures of poverty (headcount, poverty gap, squared poverty gap, Sen index and Sen-Shorrocks-Thon index) for per capita expenditure, using both the food poverty line and the total poverty line derived by the cost of basic needs method in the previous exercise.

		Food poverty line	Total poverty line
i.	Headcount index:	_____	_____
ii.	Poverty gap index:	_____	_____
iii.	Squared poverty gap index:	_____	_____
iv.	Sen index:	_____	_____
v.	Sen-Shorrocks-Thon index	_____	_____

2. Compute the headcount and poverty gap indexes for specific subgroups using the food poverty line.

		Headcount index	Poverty gap index
i.	Dhaka region:	_____	_____
ii.	Other three regions:	_____	_____
iii.	Households headed by men:	_____	_____
iv.	Households headed by women:	_____	_____
v.	Large households (>5):	_____	_____
vi.	Small households (<=5):	_____	_____

3. Repeat the above exercise using the total poverty line.

		Headcount index	Poverty gap index
i.	Dhaka region:	_____	_____
ii.	Other three regions:	_____	_____
iii.	Households headed by men:	_____	_____
iv.	Households headed by women:	_____	_____
v.	Large households (>5):	_____	_____
vi.	Small households (<=5):	_____	_____

3.3 Finding and Using `.ado` files

There is a wealth of `.ado` files on the Web, and some of them are fairly easy to locate. For example, suppose one wants to compute the Sen index of poverty. From within STATA, type **`search Sen,`** which will yield the following.

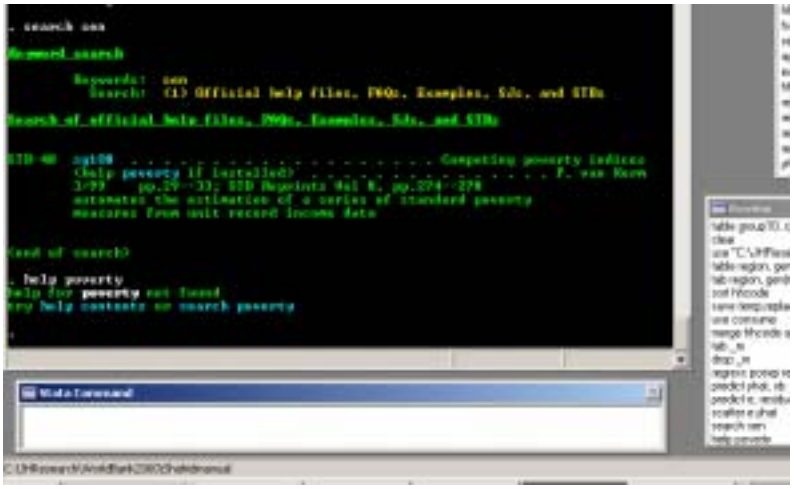

Now by double-clicking on `sg108`, you will obtain the following page, assuming that your computer is connected to the Internet.

Double-click again, this time on `click here to install`, and the relevant `.ado` file will be found, downloaded, and placed in the appropriate folder on your computer. Once this has been done successfully, you will get a screen like this one:

This file is called `poverty.ado`. To find out more about it, simply type `help poverty`. This program generates a large number of measures of poverty (but not, unfortunately, their standard errors). For a sampling of the output, try:

```
. poverty pcexp [aw=weighti], line(5000) all
```

Exercise 4 (Chapter Five):

The robustness of poverty measures is very important because if poverty measures are not as accurate as we would want, then many conclusions that we draw from poverty comparisons between groups and over time may not be warranted.

4.1 Sampling Error

For example, the fact that poverty calculations are based on a sample of households rather than the population implies that calculated measures carry a margin of error. When the standard errors of poverty measures are large, small changes in poverty may well be statistically insignificant and should not be interpreted for policy purpose.

As noted above, `FGT` also compute the standard errors of its poverty measures if option *`sd`* is specified:

```
. FGT y, Line(1000) fgt0 fgt1 sd
```

1. Now let's re-compute the headcount index and poverty gap index for Dhaka, and for the rest of the country, using the total poverty line, and compute the standard errors of the two measures as well.

		Headcount index	Poverty gap index
a) Dhaka region:	Poverty rate	_____	_____
	Standard error of poverty rate	_____	_____
b) Other three regions:	Poverty rate	_____	_____
	Standard error of poverty rate	_____	_____

2. Does the factor of standard errors change any conclusion of poverty comparison between Dhaka region and other regions?

4.2 Measurement Error

Another reason that we need to be very careful in poverty comparisons is because the data collected are measured incorrectly. This could be due to recall error on the part of respondents while answering survey questions, or because of enumerator error when the data were entered into specific formats. Let us simulate measurement error in per capita expenditure, and then investigate what effect this error has on basic poverty measures. Try the following:

```
. sum pcexp [aw=weighti]
. gen mu = r(sd)*invnorm(uniform())/10
. gen pcexp_n1 = pcexp + mu
```

Here we assume that the measurement error is a random normal variable with a standard error as big as a tenth of the standard error of observed per capita expenditure. Let us assume that the measurement error ,

mu, is additive to observed per capita expenditure. Note that, by design, this error is independent of observed per capita expenditure and of any other household or community characteristics.

1. Now re-compute the headcount ratio and poverty gap ratio using this new per capita expenditure.

		pcexp	pcexp_n1
i.	Headcount index:	_____	_____
ii.	Poverty gap index:	_____	_____

2. Are these measures different for the headcount index? For the poverty gap index?

3. Now consider the following situation. If the measurement error is correlated with a household characteristic – for example, if subsistence farmers usually underreport their consumption of own production – then will the measurement error problem be more or less severe?

4.3 Sensitivity Analysis

Apart from taking standard errors into account, it is also important to test the sensitivity of poverty measures to alternative definitions of consumption aggregates and alternative ways of setting the poverty line. For example, some non-food items are excluded from the expenditure aggregate on the basis that those items are irregular and do not reflect a household's command over resources on average. Also a 30% allowance for non-food expenditure is quite ad hoc.

(i) Create a new measure of total expenditure that includes the previously excluded irregular non-food expenditure (`expnfd2`), compute the three FGT poverty measures of per capita expenditure (`pcexp_n2`), and compare the results with those based on the original definition of expenditure (`pcexp`).

		pcexp	pcexp_n2
a.	the headcount index:	_____	_____
b.	the poverty gap index:	_____	_____

The non-food allowance can be estimated from data. Two methods have been considered (see chapter 4).

- The first finds the average non-food expenditure for households whose ***total*** expenditure is equal (or close) to the food poverty line. The non-food expenditure for this group of households must be necessities since the households are giving up part of minimum food consumption in order to buy non-food items.
- The second finds the non-food expenditure for households whose ***food*** expenditure is equal (or close) to the food poverty line. Since the second is more generous than the first, the two are usually referred

as the "lower" and the "upper" allowances and the poverty lines constructed using them are called "lower" and "upper" poverty lines, respectively.

(ii) Try the following, and then compare the results of using the two poverty lines:

```
. sum pcnfood [aw=weighti] if pcfood<foodline*1.1 & pcfood>foodline*.9
. gen line_u = foodline + r(mean)
. sum pcnfood [aw=weighti] if pcexp<foodline*1.1 & pcexp>foodline*.9
. gen line_l = foodline + r(mean)
```

	Poverty line	lower	upper
a.	the headcount index:	_____	_____
b.	the poverty gap index:	_____	_____

3. ***Challenge:*** compare poverty measures when using per adult equivalence scale expenditure (`paeexp`), with those of using per capita expenditure.

4.4 Stochastic Dominance

One may also explore the robustness of poverty comparisons by using stochastic dominance tests. The first-order stochastic dominance test compares the cumulative distribution functions of per capita expenditure. Let's compare the cumulative distributions for Dhaka with those of the rest of Bangladesh.

(i) First, generate the cumulative distribution function of Dhaka region: [Note: You may need to use the `hh.dta` file and merge it with the `consume.dta` file; you might also need to create `weighti` as the product of `weight` and `famsize`.]

```
. * Note the double equal signs to represent the identity
. keep if region == 1
. sort pcexp
. * Now create a running sum of the weighti variable
. gen cump1 = sum(weighti)
. * This normalizes cump1 so it varies between 0 and 1
. replace cump1 = cump1/cump1[_N]
. keep cump1 pcexp
. save temp, replace
```

(ii) Now generate the cumulative distribution `cump2` for the rest of Bangladesh. Keep `cump2` and `pcexp`, and append `temp.dta` by:

```
. append using temp
. label variable cump1 "Dhaka"
. label variable cump2 "other regions"
. scatter cumpl cump2 pcexp if pcexp<20000, c(l l) m(i i) title("CDFs for
Dhaka and other regions") clwidth(medthick thin)
```

(iii) Does one distribution dominate another?

(iv) If the two lines cross at least once, then we may need to test for the second-order stochastic dominance. The ***poverty deficit curve*** is the integral of the cumulative distribution up to every per capita expenditure value. After creating `cump1`, it may be obtained by

```
. gen intcump1 = sum(cump1)
. keep intcump1 pcexp
. save temp, replace
```

Create `intcump2` for the rest of Bangladesh. After combining variables and labeling them properly,

```
. label variable intcump1 "Dhaka"
. label variable intcump2 "Other regions"
. scatter intcump1 intcump2 pcexp if pcexp<20000, c(l l) m(i i)
title("Poverty Deficit Curves for Dhaka and other regions") clwidth(medthick
thin)
```

(v) Does one distribution dominate another here?

4.5 *Challenge:* Bootstrapping standard error for SST index

The bootstrapping technique can use to calculate standard errors of poverty measures, and is especially helpful in cases where the standard errors are impossible to solve analytically (e.g. with the Sen-Shorrocks-Thon index over poverty). The idea is quite simple. Repeat the calculation of the poverty measure a large number of times and each time use a new random sample drawn from the original one with replacement. For this purpose, it is necessary to use macros and loops in STATA. The following code is an example; it could be copied or typed into the do-file editor and executed.

```
set more 1
local i = 1
while `i'<=100 {
      use c:\intropov\data\final.dta , clear
      keep pcexp weighti cbnpline
      bsample _N
```

```
   SST pcexp [aw=weighti], line(5000)
      drop _all
      set obs 1
      gen sst = $S_6
      if `i' ==1  {
            save temp, replace
        }
      else {
            append using temp
            save temp, replace
       }
 local i = `i' + 1
}
sum sst
```

The above code repeats the calculation of the SST index 100 times; the `sum` command provides the standard error of these 100 estimates.

Exercise 5 (Chapter Six):

5.1 Lorenz Curve

The Lorenz curve can give a clear graphic interpretation of the Gini coefficient. Let's make the Lorenz curve of per capita expenditure distribution of rural Bangladesh.

- First, we need to calculate the cumulative shares of per capita expenditure and population: [Reminder: information on `pcexp` is to be found in `consume.dta`.]

```
. sort pcexp
. gen cumy = sum(pcexp*weight)
. gen cump = sum(weight)
. quietly replace cumy = cumy/cumy[_N]
. quietly replace cump = cump/cump[_N]
```

Second, we need to plot the cumulative share of expenditure against the cumulative share of population. It is also helpful to a 45° line (the "line of perfect equality") as a point of reference. Some of the following commands are not strictly necessary, but they do help produce a nice graph.

```
. sort pcexp
. gen equal = cump
. label variable equal "Line of Perfect Equality"
. label variable cump "Cumulative proportion of population"
. label variable cumy "Lorenz curve"
. scatter cumy equal cump, c(l l) m(i i) title("Lorenz Curve for
Bangladesh") clwidth(medthick thin) ytitle("Cumulative proportion
of income per capita")
```

- Now repeat this exercise for Dhaka region and compare its Lorenz curve with the Lorenz curve for the whole rural area. What conclusions emerge?

5.2 Inequality Measures for Rural Bangladesh

Three `.ado` programs are provided to compute the Gini coefficient, generalized entropy family and Atkinson family of inequality measures, respectively. [Note: You may need to put the `.ado` files into your working directory; or into `c:\ado\plus\g` (or `c:\ado\plus\a` in the case of the `Atkinson.ado` file)]. The `gini.ado` file is based on Deaton (1997). As in Exercise 3, you can use these programs just like other STATA commands. The syntax is:

```
. gini y [if…] [in…] [weight]
. GE y [if…] [in…] [weight], alpha(#)
. Atkinson [if…] [in…] [weight], averse(#)
```

Notes: *Alpha(#)* sets the parameter value for the generalized entropy measure that determines the sensitivity of the inequality measure to changes in the distribution. The measure is sensitive to changes at the lower end of the distribution with a parameter value close to zero, equally sensitive to changes across the distribution for the parameter equal to one (which is the Theil index), and sensitive to changes at the higher end of the distribution for higher values.
Averse(#) sets the parameter value for the Atkinson measure that measures aversion to inequality.

Let's continue using the per capita total expenditure to calculate inequality measures:

i. Compute the Gini coefficient, the Theil index and the Atkinson index with inequality aversion parameter equal to 1 for the four regions.

	Gini	Theil	Atkinson
All regions	_____	_____	_____
Dhaka region:	_____	_____	_____
Other three regions:	_____	_____	_____

ii. Now repeat the above exercise using two decile dispersion ratios and the share of consumption of poorest 25%. STATA command `xtile` is good for dividing the sample by ranking. For example, to calculate the consumption expenditure ratio between richest 20% and poorest 20%, you need to identify those two groups.

```
. xtile group = y, nq(5)
```

`xtile` will generate a new variable `group` that splits the sample into 5 groups according to the ranking of `y` (from smallest to largest, i.e., the poorest 20% will have `group==1`, while the richest 20% will have `group==5`). Similarly, to identify the poorest 25%, you need to split the sample into 4 groups.

	top 20% ÷ bottom 20%	top 10% ÷ bottom 10%	Percentage of consumption of poorest 25%
All Bangladesh	_____	_____	_____
Dhaka region	_____	_____	_____
Other regions of Bangladesh	_____	_____	_____

iii. ***Challenge*:** many inequality indexes can be decomposed by subgroups. Decompose the Theil index by region and comment on the results.

Exercise 6 (Chapter Seven):

In the previous exercises we computed poverty measures for various subgroups, such as regions, gender of head of household, household size, etc. Another way to present a poverty profile is by comparing the characteristics of the "poor" with those of the "non-poor".

6.1 Characteristics of the poor

Complete the following table, where "poor" and "non-poor" are defined by `cbnp` in Exercise 2.

	poor	non-poor
% of all households	_____	_____
% of total population	_____	_____
Average distance to paved road	_____	_____
Average distance to nearest bank	_____	_____
% of households with electricity	_____	_____
% of households with a sanitary toilet	_____	_____
Average household assets (taka)	_____	_____
Average household land holding (decimals)	_____	_____
Average household size	_____	_____
% of households headed by men	_____	_____
Average schooling of head of household (years)	_____	_____
Average age of head (years)	_____	_____
Average hh working hours on non-farm activities (p.a.)	_____	_____

6.2 More Poverty Comparison across subgroups

Calculate the headcount and poverty gap measures of poverty for the following subgroups, using `cbnpline` to define poverty.

	Headcount index	Poverty gap index
a. Household head has no education		
b. Household head has a primary education only		
c. Head had secondary or higher education		
d. Large land ownership (>0.5 ha./person)		
e. Small land ownership or landless		
f. Large asset ownership (>50000 taka)		
g. Small asset ownership (<=50000 taka)		

h. Combined with the poverty measures computed in Exercise Three, describe the most significant poverty patterns in Bangladesh?

Exercise 7 (Chapter Eight):

Develop and estimate a model that explains `log(pcexp/cbnpline)` using available data. The regressors may include demographic characteristics such as gender of head and family structure; access to public services such as distance to a paved road; household members' employment such as working hours on farm and off farm; human capital such as average education of working members; asset positions such as land holding; etc. You need to identify potentially relevant variables and the direction of their effect. Then put all those variables together, and run the regression. Report the result and discuss whether it matches your hypothesis. If not, give possible reasons.

```
. gen y = log(pcexp/cbnpline)
. reg  y age age2 workhour x1-x3 [aw=weighti]
```

where `x1-x3` are other explanatory variables that you want to include; don't feel confined to just three variables!

Note that if you want to include categorical variables, you need to convert them into dummy ("binary") variables if the ranking of categorical values do not have any meaning. For example,

```
. tab region, gen(reg)
```

will generate four variables, labeled `reg1, reg2, reg3` and `reg4`. The variable `reg1` takes on a value of 1 for Dhaka and zero otherwise, and so on. When using a set of such dummy variables in a regression, one needs to leave one of them out, to serve as a reference area. So, for instance,

```
. reg  y age age2 workhour x1-x3 reg2-reg4 [aw=weighti]
```

would include dummy variables for the regions, with Dhaka serving as the point of reference.

After the regression, it is usually a good idea to plot the residuals against the fitted values to ensure that the pattern appears sufficiently random. This could be done by adding, right after the regression command,

```
. predict yhat, xb
. predict e, residuals
. scatter e yhat
```

Index

F

G

H

I

L

M

N

O

P

Q

R

S

T

U

V