LSM- 88
FEB. 1992

Living Standards
Measurement Study
Working Paper No. 88

Poverty Comparisons

A Guide to Concepts and Methods

Martin Ravallion

Public Disclosure Authorized

LSMS Working Papers

No. 19 *The Conceptual Basis of Measures of Household Welfare and Their Implied Survey Data Requirements*

No. 20 *Statistical Experimentation for Household Surveys: Two Case Studies of Hong Kong*

No. 21 *The Collection of Price Data for the Measurement of Living Standards*

No. 22 *Household Expenditure Surveys: Some Methodological Issues*

No. 23 *Collecting Panel Data in Developing Countries: Does It Make Sense?*

No. 24 *Measuring and Analyzing Levels of Living in Developing Countries: An Annotated Questionnaire*

No. 25 *The Demand for Urban Housing in the Ivory Coast*

No. 26 *The Côte d'Ivoire Living Standards Survey: Design and Implementation*

No. 27 *The Role of Employment and Earnings in Analyzing Levels of Living: A General Methodology with Applications to Malaysia and Thailand*

No. 28 *Analysis of Household Expenditures*

No. 29 *The Distribution of Welfare in Côte d'Ivoire in 1985*

No. 30 *Quality, Quantity, and Spatial Variation of Price: Estimating Price Elasticities from Cross-Sectional Data*

No. 31 *Financing the Health Sector in Peru*

No. 32 *Informal Sector, Labor Markets, and Returns to Education in Peru*

No. 33 *Wage Determinants in Côte d'Ivoire*

No. 34 *Guidelines for Adapting the LSMS Living Standards Questionnaires to Local Conditions*

No. 35 *The Demand for Medical Care in Developing Countries: Quantity Rationing in Rural Côte d'Ivoire*

No. 36 *Labor Market Activity in Côte d'Ivoire and Peru*

No. 37 *Health Care Financing and the Demand for Medical Care*

No. 38 *Wage Determinants and School Attainment among Men in Peru*

No. 39 *The Allocation of Goods within the Household: Adults, Children, and Gender*

No. 40 *The Effects of Household and Community Characteristics on the Nutrition of Preschool Children: Evidence from Rural Côte d'Ivoire*

No. 41 *Public-Private Sector Wage Differentials in Peru, 1985–86*

No. 42 *The Distribution of Welfare in Peru in 1985–86*

No. 43 *Profits from Self-Employment: A Case Study of Côte d'Ivoire*

No. 44 *The Living Standards Survey and Price Policy Reform: A Study of Cocoa and Coffee Production in Côte d'Ivoire*

No. 45 *Measuring the Willingness to Pay for Social Services in Developing Countries*

No. 46 *Nonagricultural Family Enterprises in Côte d'Ivoire: A Descriptive Analysis*

No. 47 *The Poor during Adjustment: A Case Study of Côte d'Ivoire*

No. 48 *Confronting Poverty in Developing Countries: Definitions, Information, and Policies*

No. 49 *Sample Designs for the Living Standards Surveys in Ghana and Mauritania/Plans de sondage pour les enquêtes sur le niveau de vie au Ghana et en Mauritanie*

No. 50 *Food Subsidies: A Case Study of Price Reform in Morocco* (also in French, 50F)

No. 51 *Child Anthropometry in Côte d'Ivoire: Estimates from Two Surveys, 1985 and 1986*

No. 52 *Public-Private Sector Wage Comparisons and Moonlighting in Developing Countries: Evidence from Côte d'Ivoire and Peru*

No. 53 *Socioeconomic Determinants of Fertility in Côte d'Ivoire*

(List continues on the inside back cover)

Poverty Comparisons

A Guide to Concepts and Methods

The Living Standards Measurement Study

The Living Standards Measurement Study (LSMS) was established by the World Bank in 1980 to explore ways of improving the type and quality of household data collected by statistical offices in developing countries. Its goal is to foster increased use of household data as a basis for policy decisionmaking. Specifically, the LSMS is working to develop new methods to monitor progress in raising levels of living, to identify the consequences for households of past and proposed government policies, and to improve communications between survey statisticians, analysts, and policymakers.

The LSMS Working Paper series was started to disseminate intermediate products from the LSMS. Publications in the series include critical surveys covering different aspects of the LSMS data collection program and reports on improved methodologies for using Living Standards Survey (LSS) data. More recent publications recommend specific survey, questionnaire, and data processing designs, and demonstrate the breadth of policy analysis that can be carried out using LSS data.

LSMS Working Paper
Number 88

Poverty Comparisons

A Guide to Concepts and Methods

Martin Ravallion

The World Bank
Washington, D.C.

Manufactured in the United States of America
First printing February 1992

To present the results of the Living Standards Measurement Study with the least possible delay, the typescript of this paper has not been prepared in accordance with the procedures appropriate to formal printed texts, and the World Bank accepts no responsibility for errors.

The findings, interpretations, and conclusions expressed in this paper are entirely those of the author(s) and should not be attributed in any manner to the World Bank, to its affiliated organizations, or to members of its Board of Executive Directors or the countries they represent. The World Bank does not guarantee the accuracy of the data included in this publication and accepts no responsibility whatsoever for any consequence of their use. Any maps that accompany the text have been prepared solely for the convenience of readers; the designations and presentation of material in them do not imply the expression of any opinion whatsoever on the part of the World Bank, its affiliates, or its Board or member countries concerning the legal status of any country, territory, city, or area or of the authorities thereof or concerning the delimitation of its boundaries or its national affiliation.

The complete backlist of publications from the World Bank is shown in the annual *Index of Publications*, which contains an alphabetical title list (with full ordering information) and indexes of subjects, authors, and countries and regions. The latest edition is available free of charge from the Distribution Unit, Office of the Publisher, Department F, The World Bank, 1818 H Street, N.W., Washington, D.C. 20433, U.S.A., or from Publications, The World Bank, 66, avenue d'Iéna, 75116 Paris, France.

ISSN: 0253-4517

Martin Ravallion is principal economist in the World Bank's Human Resources Division of the Population and Human Resources Department.

Library of Congress Cataloging-in-Publication Data

Ravallion, Martin.
Poverty comparisons : a guide to concepts and methods / Martin Ravallion.
p. cm. — (LSMS working paper, ISSN 0253-4517 ; no. 88)
Includes bibliographical references.
ISBN 0-8213-2036-X
1. Cost and standard of living—Indonesia—Evaluation—Methodology. 2. Consumption (Economics)—Indonesia—Surveys—Methodology. 3. Cost and standard of living—Bangladesh—Evaluation—Methodology. 4. Consumption (Economics)—Bangladesh—Surveys—Methodology. I. Title. II. Series.
HD7055.R38 1992
339.5'7'09598—dc20 91-46743
CIP

ABSTRACT

Poverty assessments are typically clouded in conceptual and methodological uncertainties. How should living standards be assessed? Is a household survey necessary, and is it a reliable guide? Where should the poverty line be drawn, and does the choice matter? What poverty measure should be used in aggregating data on individual living standards? Does that choice matter?

This paper surveys the issues that need to be considered in answering these questions, and discusses a number of new tools of analysis which can greatly facilitate poverty comparisons, recognizing the uncertainties involved. Various applications in poverty assessment and policy evaluation for developing countries are used to show how these methods can be put into practice. Recommendations are made for future applied work.

ACKNOWLEDGEMENTS

This monograph draws heavily on results of a World Bank Research Project, "Policy Analysis and Poverty: Applicable Methods and Case Studies" (RPOs 675-04 and 675-96). I am grateful to the Bank's Research Committee for its support of the project. For their helpful comments on an earlier draft, my thanks go to Tony Atkinson, Gaurav Datt, Dominique van de Walle, Paul Glewwe, Margaret Grosh, Emmanuel Jimenez, Michael Walton, Nicholas Prescott, and Harold Alderman.

FOREWORD

It is surely undeniable that "poverty alleviation is what economic development is all about" (The World Bank Annual Report 1990, p 57). That makes it imperative that we are able to reliably assess a developing country's progress in reducing poverty, and that we have reasonable confidence about the impacts of policy initiatives and reforms on the poor. However, these tasks often pose difficult conceptual and practical problems for the analyst. This paper provides an overview of those problems, and some recommendations on how to solve them. It is aimed at the needs of economists embarking on poverty assessments or policy evaluations. The paper evolved out of Martin Ravallion's presentations at the "Economics of Poverty" Training Seminars for Bank staff, and it reflects his experience in both research and operational support on this topic. It is part of an on-going effort in the Welfare and Human Resources Division of the Population and Human Resources Department to advise World Bank Operations on how best to evaluate progress in poverty reduction in the developing world, and to help monitor that progress.

Ann O. Hamilton
Director
Population and Human Resources Department

TABLE OF CONTENTS

LIST OF TABLES AND FIGURES

1. Introduction

The most important reason for measuring poverty is probably not the need for a single number for some place and date, but rather to make a poverty comparison. This is an assessment of which of two situations has more poverty. Poverty comparisons may be either qualitative or quantitative. Examples of questions which only call for qualitative poverty comparisons are: Has poverty increased or decreased over time? Is it higher or lower in one place than another? Is there more poverty with or without some policy change? Quantitative poverty comparisons, on the other hand, call for information on how much difference there is in the amount of poverty. For some of the reasons that we want to make poverty comparisons only a qualitative assessment is essential; the choice between two policies is an example. However, in other circumstances, such as in assessing how much impact on poverty is to be expected from a specific policy option, a quantitative poverty comparison is called for.

This monograph provides a guide to the main ways that the analyst could fruitfully go about making both sorts of poverty comparisons. While some direct light is thrown on a one or two policy issues, that is incidental to the main aim, which is to expound methodological tools which can fruitfully serve policy analysis. The exposition is oriented toward the practicing applied economist who not only wants to make poverty comparisons, but also wants to be fully aware of the limitations of what is being done. Little is assumed about the economist's current knowledge of the concepts and methods of welfare and poverty analysis. But nor does the discussion shy away from possibly difficult analytical methods, as long as they show promise of facilitating sound poverty comparisons.

One theme is that, in applications, much of the energy that often goes into poverty analysis is wasted. Agonizing over where to draw some

"poverty line" is a case in point; almost always there will exist a range of possible lines over which the qualitative poverty comparison and, hence, the policy conclusion, is unaffected, and in some applications that range may be very wide indeed. The same point also holds for poverty measures, which may or may not matter to the poverty comparison.

Another theme is that too little work typically goes into assessing the robustness of poverty comparisons - both qualitative and quantitative - to changes in the underlying measurement assumptions. Much of the data we now routinely use in poverty analysis is full of errors, and that is unlikely to change. And there are unavoidable value judgements underlying measurement practice. Our policy assessments and prescriptions may or may not depend on these errors and assumptions; an important task for the poverty analyst is to find out just how confident we can really be in forming poverty comparisons.

How confident we need to be about a poverty comparison will also depend on its purpose. This illustrates a more general point - another theme of this monograph - that "measurement" and "policy" issues are often inseparable. For example, the method used to identify the poor may depend on the loss one attaches to miss-identification; the error of missing someone who is actually poor would matter more in a situation in which the poverty assessment is being used to target relief than one in which it is only being used to monitor development progress. Similarly, while it may be deemed adequate to rely solely on a simple "counting" of the poor when trying to make a summary assessment of a country's overall progress in reducing poverty, this might be a very misleading basis for choosing one policy over another. Counting the poor reveals nothing about gains and losses amongst the poor, or the extent to which the poorest are reached by a policy.

There are other aspects of the specific context of a poverty comparison which can matter in deciding precisely how that comparison

should be made. Circumstances of the set of people over which the comparison is called for - I shall call that set the "domain" - can also have bearing on measurement choices. For example, a person at a given consumption level may be deemed poor in one domain (say a generally well-off country) but not another (a poor country). However, one does have to be careful in making such arguments, and I will try to offer here a consistent framework for doing so.

Section 2 surveys the essential concepts and methods needed in applied work, while section 3 illustrates their use, drawing on my own research and policy advice in this area. Recognizing the close link between measurement and policy, some readers would no doubt prefer a more integrated approach to theory and practice than I offer here (and some have already said so). However, the set of possible policy applications is large and diverse. It is not easy to anticipate the subtleties that might arise in every possible application, or even to neatly tailor a sub-set of applications to the theory. Rather, section 2 tries to cover the full range of issues that may need to be considered - to set out the basic ideas in a reasonably generic way - leaving section 3 to illustrate the practice in selected case-studies. While facilitating the exposition, this approach may also help accommodate the quite possibly low correlation between the initial familiarity of readers with the "theory" and that for the "practice". A summary of the methodological recommendations can be found in section 4.

2. Concepts and Methods of Poverty Analysis

"Poverty" can be said to exist in a given society when one or more persons do not attain a level of material well-being deemed to constitute a reasonable minimum by the standards of that society. Saying that poverty "exists" is only the first step; for many purposes, including policy analysis, one must also say "how much" poverty exists.

The key questions for the applied economist to answer before measuring poverty are:

1. How do we assess individual well-being or "welfare"?

2. At what level of measured well-being do we say that a person is not poor?

3. How do we aggregate individual indicators of well-being into a measure of poverty?

The first two questions are sometimes referred to as the "identification problem" (which individuals are poor, and how poor are they?) while the third is called the "aggregation problem" (how much poverty is there?). Much of the theoretical literature on poverty measurement over the last 15 years or so has been concerned with the aggregation problem. However, there are also a number of difficult issues concerning the identification problem, and they are issues which can have great bearing on applied work.

2.1 Measuring "Well-Being": Conceptual Approaches

There are a number of quite different conceptual approaches to the measurement of well-being at the individual level. Approaches differ in terms of the importance the analyst attaches to the individual's own judgements about his or her well-being. They also differ in terms of the importance attached to the essentially materialist idea of "standard of living" versus less tangible but possibly no less important concepts such as "rights". I will not go into detail on the various approaches, but only try to illustrate their bearing on the key choices about

measurement which underlie poverty comparisons, such as whether to use household income or household consumption as the welfare indicator.

An important distinction is between "welfarist" and "non-welfarist" approaches (Sen 1979). The former aims to base comparisons of well-being solely on individual "utility" levels, as assessed by the individuals themselves, while the latter approach may pay little or no regard to information on utilities. This distinction is fundamental to the choices made in measuring well-being. For example, poverty comparisons in developing countries typically put a high weight on nutritional attainments. While it is clear that every individual values food consumption, one need not believe that individuals are themselves good judges of the importance of nutrition to well-being. A non-welfarist poverty comparison may thus deem that the poor are better off even if the poor do not agree.

Views differ widely on the relative merits of welfarist and non-welfarist approaches.[1] Economists have often shunned non-welfarist ideas, aiming to base assessments of well-being solely on utility information. A number of problems are encountered in doing so. If one rules out the possibility of comparing utilities across different individuals, as well as precluding non-utility information, then there can be little hope of forming consistent poverty comparisons or other judgements about social well-being.[2] Many economists are willing to admit explicit inter-personal comparisons of utility; the more contentious issue is the relevance of non-utility information, and, in particular, the extent to which one believes that individuals know what is best for themselves. There are situations where personal judgements

---

1 For example, see Sen (1979) and the comments by Ng (1981). Also see Sen (1987) and the comments therein by Ravi Kanbur and John Muellbauer. In the context of anti-hunger policies see Dreze and Sen (1989), and the comments in Ravallion (1992c).

2 This is one of the lessons to come out of the large literature on social choice theory, building on the work of Arrow (1963). See, for example, Sen (1979).

of well-being may be considered suspect, either because of miss-information, or incapacity for rational choice even with perfect information.[3]

It is common to find non-welfarist value judgements underlying policy discussions about poverty. For example, the arguments one hears in favor of "workfare", whereby recipients under anti-poverty schemes have to work to gain benefits, do not appear to be motivated by a desire to raise the utility of recipients, for that would surely be higher if one simply made a cash transfer.[4] The rationale is often not welfarist. An economist's policy advice may well fall on deaf ears if it begins with the premise that only utility information should be considered relevant.

While it would seem undeniable that utility information - the preferences of people - will often have a role of some sort in the formulation of both the objectives and the constraints in policy problems (particularly when incentives matter), welfarism is clearly not a universally acceptable principle for policy choices. Considered judgement will be needed in situations where poverty comparisons based on utility information (leaving aside for the moment the difficulties in obtaining such information) are in conflict with defensible non-welfarist comparisons. This difficult issue will come up repeatedly in poverty assessments and policy evaluations - as in other areas of policy - and I will return to it later.

The concept of "standard of living" can be either welfarist or non-welfarist. Either way, a person's standard of living is generally taken to depend solely on individual consumptions of privately supplied

---

3 For example, experimental work in psychology has suggested that "cognitive dissonance" is pervasive. This suggests that people often deal with conflicting observations about themselves in ways which generate behavior which an economist would deem to be "irrational"; for further discussion see Akerlof (1984).

4 For further discussion of this example see Besley and Coate (1988).

goods. The limitations of that approach are well recognized and access to publicly provided goods can also be included, though the problems that arise in valuing that access are often serious. Following this approach, current consumption is generally taken to be the preferred indicator of well-being in applied work, and income is only used as a proxy for consumption.

In measuring living standards, the welfarist approach typically emphasizes aggregate expenditure on all goods and services consumed, valued at appropriate prices, and including consumption from own production. (Ideally this would also include a valuation of leisure time, at an appropriate wage rate, though this is rarely feasible in practice.) By contrast a (common) non-welfarist approach emphasizes specific commodity forms of deprivation, such as inadequate food consumption (or even more narrowly, inadequate nutrition).[5] Clearly there is nothing to guarantee that someone with a total consumption which is adequate to acquire a stipulated minimum bundle of various commodities will in fact choose to consume that bundle. This is another point where there can be considerable disagreement between welfarist and non-welfarist assessments of living standards.

But "standard of living" is not the only way to think about well-being. For example, one may say that what we really want to measure is the household's "opportunity" for consumption, rather than actual consumption. To do so properly one would need data on wealth, which is rare or unreliable. But even without such data, income is clearly a better measure of opportunity for consumption than actual consumption when savings are positive. The reverse is true if savings are negative, since past savings also influence the opportunities for consumption at a given date. Thus the "opportunities approach" does not provide a fully

---

5 Following Tobin (1970), this is sometimes termed "specific egalitarianism", as distinct from "general egalitarianism". For further discussion in the context of poverty measurement see Slesnick (1990).

compelling argument for preferring income to consumption as the welfare indicator for all households.

Non-welfarist ideas of "rights" can also have bearing on choices made in measuring well-being. This approach says that we are concerned about poverty because we are concerned about the attainment of the right to participate in a society, and this depends crucially on income, particularly cash income. By this approach, one not only prefers income, one may want to weight different components of income differently than is usually the case; for example, income received from the state but with some social stigma may add little to one's ability to participate in society (Atkinson 1991).

The standard of living approach has been more popular in development literature (it is, for example, the approach implicit in the World Development Report 1990; see World Bank 1990b) and generally dictates a preference for consumption as the welfare indicator. Ideas such as "opportunities" and "rights" seem to have carried relatively more weight in the developed country literature, particularly in Europe, and they are generally seen to indicate a preference for income (Atkinson 1991), though subject to the exceptions noted. I suspect that the greater importance that has been attached to the standard of living approach in developing country literature reflects the greater importance attached to specific forms of commodity deprivation, especially food insecurity. That emphasis is quite defensible, from both welfarist and non-welfarist points of view. The rest of this monograph will focus solely on the standard of living approach, though I will consider both "welfarist" and "non-welfarist" perspectives.

2.2 Household Surveys

Household surveys are the single most important source of data for making poverty comparisons; indeed, they are the only data source which can tell us directly about the distribution of living standards in a

society, such as how many households do not attain some consumption level. However, a lot of care must go into setting up and interpreting such data. This section will survey the issues that the analyst should be aware of.

The household surveys found in practice can be classified along three dimensions:

i) The unit of observation: This can be the household itself or the individuals within the household. In practice most household surveys include data on individuals within the household, though this rarely includes their consumptions, which are typically aggregated to the household level; examples include India's National Sample Surveys (NSS), Indonesia's National Socio-Economic Surveys (SUSENAS), the World Bank's Living Standards Measurement Surveys (LSMS) and the surveys done as part of the project on the Social Dimensions of Adjustment in Sub-Saharan Africa, based at the Bank. An example of a survey which collected individual food consumption data is the survey of rural households in the Philippines done by the International Food Policy Research Institute (Bouis and Haddad 1990).

ii) The number of observations over time: A single cross-section, based on one or two interviews, is the most common. A few longitudinal surveys (also called panel data sets) have been done in developing countries, in which the same household is re-surveyed over an extended period. An example is the Village Level Survey over ten years by the International Crops Research Institute for the Semi-Arid Tropics in India. There are also a few surveys which have collected individual consumption data on a longitudinal basis (such as the aforementioned IFPRI survey for the Philippines). However, these rarely have wide coverage within a country, as they are expensive surveys to collect, and they may suffer problems of "panel ageing", whereby the initial sample becomes increasingly un-representative over time. A modified version of longitudinal sampling has been used in some LSMS surveys, whereby half

of each year's sample is re-surveyed the following year. This cuts the cost of forming a panel data set, while retaining some of the advantages.

iii) The principal living standard indicator collected: The most common indicators used in practice are based on household consumption expenditure and household income. Some surveys collect both variables (such as Indonesia's SUSENAS and the Bank's LSMS), but others specialize (for example, India's NSS does not include incomes, while most of the household surveys available for Latin America do not include consumption).

The most common survey uses the household as the unit of observation, is for a single cross-section, and collects either consumption or income data. This form of survey is cheaper per household surveyed than the alternatives, and this allows a larger sample than with longitudinal and or individual based survey. A large sample of household level data gives greater accuracy in estimating certain population parameters, such as average consumption per capita, but, of course, looses accuracy in estimating other variables, such as the number of under-nourished children in a population. It should not be presumed that the large household consumption survey is more cost-effective for all purposes than alternatives using smaller samples of individual data.

The following are the main problems for the practitioner to be aware of when interpreting household consumption or income data from a household survey.

Survey Design

Even very large samples may give biased estimates for poverty measurement if the survey is not random.[6] A random sample requires

---

6 For a good introduction to sampling methods see Levy and Lemeshow (1991). Also see UN (1989).

that each person in the population has an equal chance of being selected for the sample. This guarantees statistical independence, the assumption which underlies most of the results used routinely in making statistical *inferences* about population parameters from sample surveys. Sample statistics, such as the proportion of people who do not obtain a particular consumption level, are un-biased and have well-defined distributions for random samples.

However, the poor need not be properly represented in sample surveys; for example they may be harder to interview because they live in remote areas or are itinerant. Indeed, a household survey will 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 frame was not intended to span the entire population; examples include the "labor force surveys" (widely used for poverty assessments in Latin America) for which the sample frame is typically restricted to the "economically active population" which precludes certain sub-groups of the poor. Key questions to ask about the survey you are using are: Does the sample frame (the initial listing of the population from which the sample was drawn) span the entire population? Is there likely to be a response bias, in that the likelihood of cooperating with the interviewer is not random?

There are various methods of sampling which can help achieve a more cost-effective survey than would be possible with simple random sampling. 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; for example, one can over-sample certain regions where the poor are thought to be concentrated.[7] Cluster sampling can help cut

---

7 Though it is important to know the sampling rates, so as to correctly weight the sub-samples when forming a population estimate; see Section 2.4.

survey costs, particularly when (as is common in developing countries) a wide coverage sample frame is not available. By this method, one first randomly samples clusters of households, such as blocks in a city, and then samples randomly within the selected clusters. However, the lower cost of cluster sampling must be weighed against the fact that clusters tend to be more internally homogeneous than in a random sample, which adds to the imprecision of sample estimates. If cluster sampling has been used in your survey, it is often important to know how this was done; for example, if only one cluster was picked in each of the regions which constitute your poverty profile then the profile may be quite misleading.

Goods Coverage and Valuation

The coverage of goods and income sources in the survey should be comprehensive, covering both food and non-food goods, and all income sources. Consumption should cover all monetary expenditures on goods and services consumed plus the monetary value of all consumption from income in kind, such as food produced on the family farm and the value of owner-occupied housing. Similarly, the income definition should include income in kind.[8] Whenever prices are unknown, or are an unreliable guide to opportunity costs, valuation problems will arise. Such problems are common for consumption and income derived from the household's own production, such as farm output.[9] The valuation of benefits from public services is also notoriously difficult and few survey analysts have attempted to do so.

A common problem facing the welfare analyst using a household survey is that the survey may not be properly integrated, in that

---

8 For further discussion of consumption and income definitions in household surveys see UN (1989).

9 For a sobering example of the measurement errors that can arise from an imperfect acccounting of consumption from own-farm production see Gautam (1991).

categories do not match in relevant ways across different segments of the survey. For example, to evaluate the welfare effects of a change in food staple prices in a food producing developing country it is not enough to know the budget shares of consumption at the household level - one must also know household food production. Whether a household gains or loses from a change in the price of food depends on consumption net of production. (Section 3.7 gives examples.) Furthermore, unlike developed countries, income from farm production is an important income source for many of the poor and near poor in developing countries. However, it is quite common that household surveys for developing countries either do not include data on farm production, or they do not use the same commodity categorization in the consumption and production schedules. This often makes those surveys virtually useless for analyzing an important class of development policy problems.

Variability and the Time Period of Measurement

Inter-temporal variability has implications for a number the choices made in measurement using survey data. One is the choice between income and consumption based measures.

Most analysts using household data for developing countries in making welfare comparisons have preferred current consumption to income as the indicator of living standards. Variability is probably the main reason. Incomes of the poor often vary over time in fairly predictable ways (and sometimes in quite un-predictable ways); this is particularly true in underdeveloped rural economies depending on rain-fed agriculture. Typically, there are consumption smoothing and insurance opportunities available to the poor, such as through saving and community based risk-sharing. This observation has two distinct implications for welfare measurement: i) current consumption will almost certainly be a better indicator than current income of <u>current</u> standard of living, and ii) current consumption may then also be a good indicator

of long-term well-being, as it will reveal information about incomes at other dates, in the past and future.

However, a number of factors can make current consumption a "noisy" indicator of long-term well-being. Even with ideal smoothing, consumption will still (as a rule) vary over the life-cycle. Thus two households with different life-time wealth - one "young", the other "old" - may happen to have the same consumption at the survey date.[10] This is probably less of a problem in traditional societies where the extended family is still the norm. However, there are other sources of noise in the relationship between current consumption and long-term standard of living. 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 borrowing options than the non-poor (also suggesting that life-time wealth is not the only parameter of life-time welfare). While consumption smoothing and risk-sharing arrangements clearly do exist, their performance from the point of view of the poor is a moot point.[11] Furthermore, even if one accepts that current consumption is a more precise estimator of long-term well-being than current income for a given household, that does not mean it will be the better ordinal indictor of who is poor in terms of typical long-term living standard. That also depends on how the various living standards indicators rank different households; one cross-sectional indicator may vary less around long-term living standard than another but cause more re-ranking and, hence, perform less well in identifying the chronically poor. (This last issue is taken up in more detail in section 3.2.)

---

10 See Blundell and Preston (1991) for an analysis of this issue in a an inter-temporal model of consumption.

11 For a sample of relevant empirical work for developing countries see Deaton (1991), Bhargava and Ravallion (1991), Townsend (1990) and Ravallion and Chaudhuri (1991).

Some of the aforementioned problems of welfare measurement have implications for survey design. Since many of the rural poor face marked seasonality, income over a whole year will better reflect living standards for agricultural households than one quarter say. However, interviewee recall is imperfect, and can rarely be relied upon to span the frequency of income variation. Thus consumption will often be a better guide, as discussed above. Careful survey design can also enhance precision in estimating consumption. Better estimates can generally be obtained by adapting the period of recall to the frequency at which the type of good is purchased; recall for one week may be fine for food,[12] while a three month period is probably more appropriate for clothing, say. With panel data one can enhance precision in estimating living standards by averaging the multiple consumption or income observations over time.[13] But such surveys are more costly.

Comparisons Across Households at Similar Consumption Levels

Household size and demographic composition vary, as do prices and access to publicly supplied goods. These generate differences in well-being at given household expenditures. There are various welfarist approaches based on demand analysis, including equivalence scales, cost-of-living indices, and equivalent income measures.[14] 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 function is derived which is consistent with observed demand behavior, relating consumption to prices, incomes,

---

12 Seasonality in food consumption may warrant a longer recall period than this, though recall error will then increase, so there may be little or no final gain in precision.

13 See, for example, Ravallion (1988), Gaiha and Deolalikar (1989), World Bank (1990, Chapter 2), Lanjouw and Stern (1991), Chaudhuri and Ravallion (1991).

14 There are a number of good expositions of these topics; see, for example, see Deaton and Muellbauer (1980) and Deaton (1980).

household size, demographic composition, and other relevant variables which shift tastes.[15] The next section will discuss the use of "equivalence scales", while an example of the use of "equivalent incomes" is discussed in section 3.7.

In all such behavioral welfare measures, the basic problem to be aware of is that a given set of revealed preferences over goods may be consistent with infinitely many reasonable ways of making inter-personal welfare comparisons; it is a big step to assume that a particular utility function which supports observed behavior as an optimum is also the one which should be used in measuring well-being.[16] For example, I would be surprised if the extra satisfaction that parents derive from a new baby is fully evident in their consumption behavior.

If one cannot come up with a convincing way of measuring differences in welfare at given consumption levels then one should do poverty comparisons separately for each of the different groups between which welfare is thought to vary at given consumptions (such as different household sizes). Sections 2.4 and 2.6 return to this point.

There are other data which can also help. Ideally we should not have to rely solely on a household level survey in making interpersonal comparisons of welfare. A separate community survey (done at the same time as the interviews, and possibly by the same interviewers) can provide useful supplementary data on the local prices of a range of

---

15 This method assumes that the parameters of the demand system satisfy the theoretical conditions of utility maximization (see,for example, Varian 1984). The utility function is derived from the estimated demand model either as an explicit functional form (see, for example, King 1983) or by numerical methods (Vartia 1983). For an exposition of the theory as it applies to the construction of equivalence scales see Deaton and Muellbauer (1980, 1986).

16 For further discussion see Pollak and Wales (1979), Deaton and Muellbauer (1986), and Fisher (1987).

goods and local public services.[17] By matching this to the household level data one can greatly improve the accuracy and coverage of household welfare assessments. This has become common practice in the World Bank's LSMS surveys.

A useful strand of recent research on living standards measurement has been concerned with the comparison of how different indicators at the individual or household level identify different individuals as poor.[18] Much can be learnt from individual and/or longitudinal surveys about the properties of the more common forms of (less costly) single cross-sectional household survey. Section 3.2 gives an example.

2.3 Some Alternative Measures

Consumption per Equivalent Adult

I have already looked at the arguments for and against preferring consumption to income as the living standards indicator, and section 3.2 will return to that issue. I now want to comment further on the use of equivalence scales in comparing living standards across households.

Households differ in size and composition, and so a simple comparison of aggregate household consumption could be quite misleading about the well-being of individual members of a given household. Most analysts recognize this problem, and use some form of normalization, such as "consumption per equivalent adult male". For a household of any

---

17 The recall of interviewees on public services and the prices implicit in their reported quantities consumed and expenditures on various goods can be used for these purposes. However, there are a number of problems. Knowledge about local public services depends on usage which may be a biased indicator of actual availability. Retrieving prices (strictly "unit values") from quantities and expenditures can also give biased results when quality varies in unknown ways, though there are remedial methods of dealing with this problem (Deaton 1988). However, it remains that prices for non-food goods cannot be obtained this way (as the data rarely allow meaningful comparisons, so only expenditures are obtained). A good community schedule can deal quite well with all these problems.

18 Examples include Anand and Harris (1990), Glewwe and van der Gaag (1990), Haddad and Kanbur (1990), Lanjouw and Stern (1991), and Chaudhuri and Ravallion (1991).

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) which that household is deemed to be equivalent to. The key question is: "equivalent" in what sense?

In the practice of assigning equivalence scales, the answer is typically based on observed consumption behavior from household surveys. In essence, one looks at how aggregate household consumptions of various goods during some survey period tends to vary with household size and composition (as well as prices and incomes) over the cross-section of households surveyed. For example, by one common method, a demand model is constructed in which the budget share devoted to food consumption of each household is regressed on the log of total consumption per person and the numbers of persons in various demographic categories living in the household (this is the Working-Leser form of the Engel curve). Under certain conditions the food-share can be interpreted as a welfare indicator (as discussed further below). By fixing some reference welfare level and, hence, food-share one can then use the regression equation to calculate the difference in total consumption per person which would be needed to exactly compensate one household for its different composition to that of another household (see, for example, Deaton and Muellbauer 1986).

The result of such methods is that most equivalence scales tend to assign an adult male equivalence less than one to adult females and children; females and children tend to consume less of most goods than do adult males. It is often assumed that this reflects a difference in "needs"; that women and children need less consumption to achieve the same level of well-being as men.

Is this practice to be recommended? There are a number of problems. The example discussed above based on an estimated Engel curve assumes that the food-share is a valid welfare indictor; the possible problems in that assumption are discussed below. Of course, if one is

happy with that assumption then one need not bother with estimating equivalence scales for welfare and poverty measurement - the food-share is sufficient data.

There are also some theoretical problems in this practice. One was mentioned above, namely that the welfare interpretation of observed consumption behavior is clouded by the fact that there will often be multiple utility functions (indeed infinitely many) which generate the same behavior. Relevant parameters of well-being will not then be identifiable from that behavior. Another issue is that child costs can also be financed by drawing on savings rather than consumption, so that the effect on consumption may occur at a later date than the survey (the children may even have grown up).[19] Purely static observations of consumption and household demographic characteristics can thus be a misleading guide in forming equivalence scales.

The welfare interpretation of equivalence scales constructed from consumption behavior also depends on the view one takes about how consumption allocations are made within the household. The interpretation of the empirical evidence on which equivalence scales are based may be quite different if the data are assumed to be generated by a male dictatorship (at one extreme) rather than the maximization of a function of the well-being of all household members. The most plausible model would seem to be one of bargaining within the household, in which intra-household allocations will reflect the outside options of household members.[20] The equivalence scale derived from consumption behavior can then be taken to embody two distinct aspects of distribution within the household: real differences in "needs" between certain age and gender groups, and inequalities in outside options or

---

19 For an analysis of the implications of inter-temporal consumption behavior for the estimation of equivalence scales see Pashardes (1991).

20 For a survey of such models see McElroy (1990). Also see Sen (1984 Chapter 15), Schultz (1990), Thomas (1990), Haddad and Hoddintott (1991).

"bargaining power". While the analyst and policy maker would rightly want to incorporate the first into household welfare comparisons, one would be loathe to incorporate the second, as this would perpetuate and even reinforce an existing welfare inequality.

The potential policy implications of this measurement problem can be illustrated with a simple example, using the hypothetical data given in Table 1. There are four persons, living in two households. Household A has one adult male, one adult female and two children, while B comprises a single adult male. Individual consumptions are given in Table 1. In terms of those consumptions the three poorest persons are in household A. To make the example sharper, I shall assume that this is also true when consumptions are normalized for differences in "needs". The government can make a transfer to the household which is deemed to be the poorest, but it cannot observe distribution <u>within</u> any household; all the government knows is aggregate household consumption, and household composition.

Table 1: Consumptions Within Two Hypothetical Households

Household	Individual consumptions of				Household consumptions	
	Male adult	Female adult	First child	Second child	per person	per equivalent male adult
A	40	20	10	10	20	40
B	30	--	--	--	30	30

Which of the two households, A and B, should be first to receive help? As long as at least some of it benefits women and children, the answer is clearly household "A". But to know this you would have to know individual consumptions. In terms of household consumption per person, which is known, the answer is also A. Using this equivalence

scale - which weights all persons equally - at least some of the benefits will go to the three poorest persons. However, consider instead a household scale which assigns weights proportionally to actual consumption levels. (As might be obtained by running a regression on a sample of households with similar consumption levels and compositions to those in Table 1.) The equivalence scale would be 0.5 for an adult female, and 0.25 for each child. There are thus two equivalent adult males in household A, which then has a consumption per equivalent adult male which is more than that of household B. B will receive help first, and none of it is likely to go to the poorest 60% of the population.

Of course, this is only an example, and one based on (possibly) quite extreme inequality within household A. However the example is adequate to demonstrate two key points:

i) While observable consumption behavior is important data, assumptions about unobservables will be required.

ii) Seemingly innocuous assumptions made when making inter-household comparisons of well-being in empirical work can have considerable bearing on policy choices.

Food-Share

As with any non-luxury good, the budget share devoted to food tends to decrease with total real consumption expenditure. This observation, sometimes called "Engel's Law", has often been invoked to justify using the non-food budget share as an indicator of living standards.

There are a number of problems with this indicator. For one thing, the relationship between the food budget share and consumption will generally differ across households (due to differences in the relative prices they face, demographic differences, or differences in preferences). This creates noise in the indicator. Also the income elasticity of demand for food can be very close to unity for poor

households, and then the food share can be a quite un-reliable indicator. Section 3.2 discusses an example.

However, food share data can sometimes provide a useful supplementary test, particularly if one is worried about the quality of other data. The worry may stem from either the survey data or the price deflator. For example, Ravallion and Huppi (1991) find that their food share data gives the same qualitative conclusions in comparing poverty over time and sectors in Indonesia as did consumption and income data, adding strength to the paper's conclusion that poverty in Indonesia continued to decline during the study period 1984-1987; see section 3.4 for further discussion.

Nutritional Indicators

As both terms are commonly understood, "undernutrition" is a distinct concept to "poverty". However, the difference is in the definition of the individual measure of "well-being" used - nutrient intakes (notably food energy, but also micronutrients) versus a broader concept of "consumption" which includes other attributes of food besides their nutritional value, and non-food consumption. Thus, in a formal sense, one can view undernutrition as "food energy poverty", and measure it in a similar way.

There are arguments for and against using nutrient intakes as an indicator of well-being in low-income countries. As with food-share, a practical advantage in countries with high rates of inflation, or inadequate price data, is that distributional data on food energy intakes do not need to be adjusted for inflation.[21] However, against this, nutrition only captures one aspect of well-being. In low-income countries, food staple consumption will have a high weight in any

---

21 Though that does not mean that food energy intakes are unaffected by inflation or changes in relative prices; but these are not things we need worry about in measuring changes in undernutrition.

demand-consistent welfare indicator, but it will rarely have a weight of one.

Again it may be argued that consumption behavior is not a good enough guide for welfare measurement; the weight people attach to food, and nutrient intakes in particular, may be considered "too low for their own good". My own view is that, while one should be wary of purely welfarist arguments which assume that people are always the best judges of their own welfare, one should be equally suspicious of any measure of living standards which ignores demand behavior. Given the obvious uncertainties on this issue, the only sensible solution seems to be to monitor selected "non-welfarist" indicators side by side with "welfarist" ones. Only if the two types of measures disagree on the poverty comparison need one delve further into the issue. When one has to do so, a convincing non-welfarist assessment should, in my view, identify plausible reasons why revealed preference is inconsistent with well-being. Are there reasons why consumption behavior is misguided, such as due to the intra-household inequalities discussed above in the context of equivalence scales? Is it an issue of imperfect information (with implications for education policies)? Or is it a more fundamental problem, such as irrationality (due, for example, to cognitive dissonance) or incapacity for rational choice (such as due to simply being too young to know what is good for you, and not having someone else to make a sound choice).

The above comments also apply to anthropometric measures of the nutritional status of children, such as weight-for-age or weight-for-height. Their appeal stems in part from the uncertainties in the measurement of individual nutritional requirements, though not dissimilar uncertainties pervade the choice of anthropometic standards.

However, there is one further point about these measures: by some accounts (including some nutritionists) 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 (Beaton 1983). That is clearly a serious food related deprivation for any child. Again, one should, in my view, be wary of overly narrow conceptualizations of the meaning of individual "well-being" when making poverty comparisons.

Anthropological Methods

While it is clearly not a feasible method for national level poverty comparisons over time (say), close observation at the household level over an extended period can provide useful supplementary information on living standards. For example, Lanjouw and Stern (1991) used subjective assessments of poverty in a north Indian village, based on the observations of resident investigators over one year. This involved classifying households into seven groups (very poor, poor, modest, secure, prosperous, rich, very rich) on the basis of observations and discussions with villagers over that year.

An issue of concern about this method is 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 we find that being a landless agricultural laborer in their surveyed village is virtually a sufficient condition for being deemed "apparently poor" by their anthropological method - 99% of such households are deemed poor by this characteristic, though this is only so for 54% when their measure of permanent income is used, based on averaging current incomes over four interviews spanning 25 years. 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 type of data is telling us the most about the reality of poverty.

2.4 Poverty Lines

Poverty measurement generally assumes that there exist pre-determined and well-defined standards of consumption - called "poverty lines" - which must be reached if a person is not to be deemed "poor". It is undeniable that there exist levels of consumptions of various goods (food, clothing and shelter) below which survival beyond short periods is threatened, though it is less clear what these levels exactly are for any given individual. Furthermore, in most societies - including some of the poorest - the notion of what constitutes "poverty" goes beyond the attainment of the absolute minimum needed for survival. Poverty lines exist, but views differ on their location.

There are numerous approaches to resolving this uncertainty.[22] I will provide a critical overview of the main alternatives found in practice, focusing particularly on their relevance to developing countries.

Absolute Poverty Lines

Much of the literature and policy discussion in developing countries has concentrated on absolute poverty. This should not be defined, in my view, as the use of an especially stringent ("survival") poverty line. Rather, an absolute poverty line is one which is fixed in terms of the living standards indicator being used, and fixed over the entire domain of the poverty comparison. Thus absolute poverty comparisons will deem two persons at the same real consumption level to both be either "poor" or "not poor" irrespective of the time or place

---

22 For surveys (though more from a developed country perspective) see Hagenaars and de Vos (1988), and Hagenaars and van Praag (1985).

being considered, or with or without some policy change, within the relevant domain.

The failure of many discussions to recognize the specificity of poverty comparisons to their domain can be a source of confusion. For example, when trying to make a global comparison of absolute poverty in terms of consumption, there is (in my view) a compelling case for using the same real consumption level as the poverty line for all countries. That will probably entail using a poverty line in a rich country which is low by the standards of that country. But the domain of that particular poverty comparison goes well beyond the borders of one country. If, however, one was trying to form a poverty profile for one country only, the choice of an absolute poverty line should be appropriate to that country. Judgements of what constitutes a reasonable absolute poverty comparison must first specify the domain, and recognize that the answer may change if the domain changes.

The most common approach in defining an absolute poverty line is to estimate the cost of a bundle of goods deemed to assure that basic consumption needs are met in the specific domain of the poverty comparison.[23] The difficulty is in identifying what constitutes "basic needs". For developing countries, the most important component of a basic needs poverty line is generally the food expenditure necessary to attain some recommended food energy intake. This is then augmented by a modest allowance for non-food goods.

The first problem is choosing the food energy requirement; this can vary across individuals and over time for a given individual. An assumption must be made about activity levels which determine energy requirements beyond those needed to maintain the human body's metabolic rate at rest. Activity levels are, however, "endogenous" socio-economic

---

23 The basic needs approach to defining poverty lines goes back to the work of Rowntree, in his study of poverty in York, England, at the turn of the nineteenth century. For a description of Rowntree's study see Atkinson (1975, Chapter 10).

variables rather than "exogenous" physiological ones. Another problem is that the minimum cost of the stipulated number of calories may be a good deal less than the expenditure level at which the poor typically attain that calorie level. Attaining adequate nutrition is not the sole motive for human behavior, even for most of the poor; nor is it the sole motive in food consumption.

The second problem is making an allowance for non-food consumption. One method - which I shall term the "food energy method" - for implementing the basic needs approach proceeds by first fixing a food energy intake cut-off in calories, and then finding the consumption expenditure or income level at which a person typically attains that food energy intake. This can be estimated from a regression of calorie intake against consumption expenditures or income.[24] The method automatically includes an allowance for non-food consumption, as long as one locates the total consumption expenditure at which a person typically attains the caloric requirement. It also has the appeal that it yields a poverty line which is consistent with local tastes, as well as prices.

A variation on this method is to first find the minimum cost of a food bundle which achieves the stipulated energy intake level, and then divide this by the share of food in total expenditure of some group of households deemed likely to be poor. This is sometimes called the Orshansky method, following Orshansky (1965) who used it to measure poverty in the USA.

The food energy method is fine for setting a single poverty line, but one should be careful in applying it separately to each region, sector or date in the poverty comparisons being made. If one is comparing living standards in terms of household consumption per capita then comparisons of absolute poverty across regions, sectors or dates can be misleading unless the poverty line has constant purchasing power

---

24 For an example see Greer and Thorbecke (1986).

(based on a cost-of-living index appropriate to the poor). However, the above methods are quite unlikely to generate poverty lines which are constant in terms of real consumption or income. The reason is that the relationship between food energy intake and consumption or income is not going to be the same across regions/sectors/dates, but will shift according to differences in tastes, activity levels, relative prices, publicly provided goods or other variables. And there is nothing in this methodology to guarantee that these differences are ones which would be considered relevant to poverty comparisons. In the Orshansky method, differences will also arise simply because of differences in average real consumption or income across groups or dates; those with a higher mean will tend to have a lower food share which will thus lead one to use a higher poverty line.

The differences can be large enough to cause a rank reversal in measured poverty levels across sectors or regions of an economy. Section 3.3 gives an example. This can be worrying when there is mobility across the groups being considered in the poverty profile, such as migration from rural to urban areas, as discussed in section 3.3.

In short, when applied separately to each of the situations being compared, some common methods of setting poverty lines risk confounding the ideas of "absolute" and "relative" poverty, and in ways which do not make clear whether we are observing changes in absolute poverty or changes in relative poverty. Furthermore, it is not even clear that the way in which relative poverty considerations are entering into the poverty comparison is sensible.

There are other reasons why one might set different poverty lines for different sub-groups of a population when making absolute poverty comparisons. One way of dealing with the possibility that the living standards indicator does not properly reflect differences in well-being at a given consumption level is to set different poverty lines. This may be considered easier than revising the living standards indicator to

better reflect the differences in needs. However, this may also be a rather restrictive way of dealing with differences in needs, since it need not yield meaningful comparisons across different needs groups for those below the poverty line; comparability is only assured at the poverty line. One might try to get around this problem by additively adjusting the living standards indicators for the difference between the poverty lines, though there is nothing to guarantee that the differences in needs among all of the poor is properly accounted for by adding a constant at all levels.

Given that there is a degree of arbitrariness in the basic needs approach, an alternative way of setting a poverty line is to first identify as "the poor" the poorest x% of the population at some base date or place, and use the corresponding consumption or income level for this percentile as the poverty line for comparisons with other dates or places. Again, this is only an absolute poverty line if it is fixed in terms of the living standards indicator over the domain of the poverty comparison (the value of x will then, of course, vary).

The most important point in the above discussion is that, recognizing that a certain amount of arbitrariness is unavoidable in defining any poverty line in practice, one should be particularly careful about how the choices made affect the poverty comparisons, for these are generally what matter most to the policy implications.

Relative Poverty Lines

Another difference between the developing country and developed country literatures is that absolute poverty considerations have dominated the former, while relative poverty has been more important in

the latter. For example, many studies for developed countries have used a poverty line which is set at about 50% of the national mean.[25]

Is there a compelling case for using poverty lines set at a constant proportion of the mean in developing countries? I shall discuss poverty measures in greater detail in the following section, but for now we need only note that almost all measures of poverty are homogeneous of degree zero between the mean and the poverty line. We can write such a poverty measure in the following generic form:

$$(1) \qquad P = P(z/\mu, \mathbf{L})$$

where z is the poverty line, μ is the mean of the distribution on which poverty is measured, and **L** is a list of parameters fully describing the Lorenz curve of that distribution, which summarizes all relevant information about relative inequalities. With a poverty line set at $z=k.\mu$, where k is some constant (such as 0.5, as often used in the European studies cited by Atkinson 1991), the measure of poverty becomes P(k, **L**), and depends <u>solely</u> on the parameters of the Lorenz curve. If all incomes increase by the same proportion then P(k, **L**) would remain totally unchanged - there would be no change in relative inequalities and so **L** would not change. And the poverty line would simply increase by the same proportion.

It might be argued that P(k, **L**) is still a good measure of "relative poverty", to the extent that what one is really trying to capture in this concept is the amount of inequality in the distribution, which can be thought of as depending solely on the Lorenz curve. However, then we should ask whether or not a ranking of distributions in terms of P(k, **L**) will preserve their ranking in terms of an appropriate

---

25 Atkinson (1991) shows how poverty comparisons across countries in Europe are affected by this choice; there is substantial re-ranking when one compares poverty measures based on a constant proportion of each country's mean income with those obtained using the same proportion applied to a constant mean across <u>all</u> countries.

measure of inequality. Any measure of inequality should respect the principle that whenever income is transferred between two persons, inequality will have decreased (increased) if the donor had a higher (lower) income than the recipient (Atkinson 1970, 1975). One can readily construct examples whereby distribution A Lorenz dominates B - so that A has less inequality than B for any well-behaved measure of inequality - and yet P(k, L) is higher for the A distribution.[26] And such examples are also possible when the transfers are only made amongst the poor. Thus P(k, L) is not only independent of the mean, it need not be consistent with reasonable normative judgements about relative poverty.

Part of the problem here arises from the assumption that the poverty line is a constant proportion of the mean. This implies that the elasticity of the poverty line w.r.t the mean is unity. A survey of poverty lines across 36 countries, both developing and industrialized, was done for World Development Report 1990, and the results were compared to average private consumptions in those countries (Ravallion, Datt and van de Walle, 1991). The result is plotted in Figure 1. The elasticity of the poverty line w.r.t. the mean is increasing in the

---

26 For example, consider the simplest poverty measure, the proportion of people deemed to be poor (the "head-count index"). Suppose we want to compare poverty in two states A and B - they might be different countries, different regions, different socio-economic groups, or different dates for the same population. Let the Lorenz curves for A and B be $L_A(H_A)$ and $L_B(H_B)$ respectively. The head-count indices are then H_A and H_B where these are given by:

$$L'_A(H_A) = k = L'_B(H_B)$$

(Noting that $L'(p)\mu=x$ which is the inverse of the cumulative frequency distribution, giving the proportion of the population below any point x.) Now suppose than the mean of A exceeds that of B, and that the distribution in A Lorenz dominates that of B (i.e., lower inequality in A than B). Even when these conditions hold, one can always find valid Lorenz curves such that $H_A > H_B$. The estimated poverty measure could indicate higher poverty in A than B, though inequality and absolute poverty are unambiguously lower in A than B.

Figure 1: Poverty Lines for 36 Countries

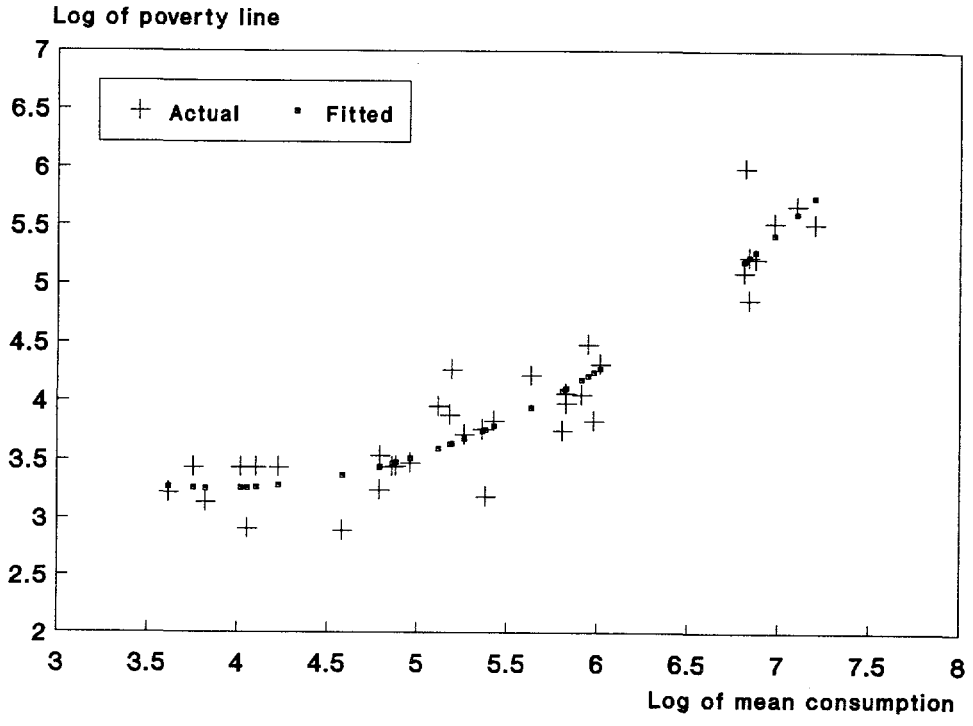

Note: The poverty lines are plotted against private consumption per capita from national accounts, both at purchasing power parity around 1985. Fitted values based on a regression of log poverty line against a quadratic function of log mean consumption.

Source: See Ravallion, Datt and van de Walle 1991; data available from the author.

mean. At the mean of the country means the elasticity is .66. However, at the mean consumption of India, the elasticity is very much lower, at 0.15. At the mean consumption level in the poorest country it drops to 0.07. But, amongst the industrialized countries, the elasticity is about unity.

In short: the cross-country comparison does suggest that real poverty lines will tend to increase with growth, but they will do so slowly for the poorest countries. Notions of "absolute poverty" - whereby the poverty line does not vary with overall living standards -

appear to be relevant to low income countries, while "relative poverty" is of more relevance to high income countries. Furthermore, the proportionality assumption often made in the developed country literature appears to be quite reasonable for the advanced industrialized countries, though the measure obtained is very difficult to interpret in terms of conventional concepts of inequality and poverty.

Subjective Poverty Lines

This approach explicitly recognizes that poverty lines are inherently subjective judgements people make about what constitutes a socially acceptable minimum standard of living in a particular society. Just as the previous section showed how different countries tend to use different poverty lines, and that richer countries tend to have higher poverty lines, so too with individuals.

This approach is often based on survey responses to a question such as the following (paraphrased from Kapteyn et al 1988):

> "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 answer tends to be an increasing function of actual income. Furthermore the studies that have included this question (mainly for Europe) have tended to find a relationship as depicted in Figure 2 (Kapteyn et al, 1988). The point z in the figure is an obvious candidate for a poverty line; people with income above z tend to feel that their income is adequate, while those below z tend to feel that it is not.

This is an interesting approach to setting a poverty line, though I know of no attempts to apply it in a developing country. I would not be surprised if it gave higher poverty lines than the basic needs

Figure 2: The Subjective Poverty Line

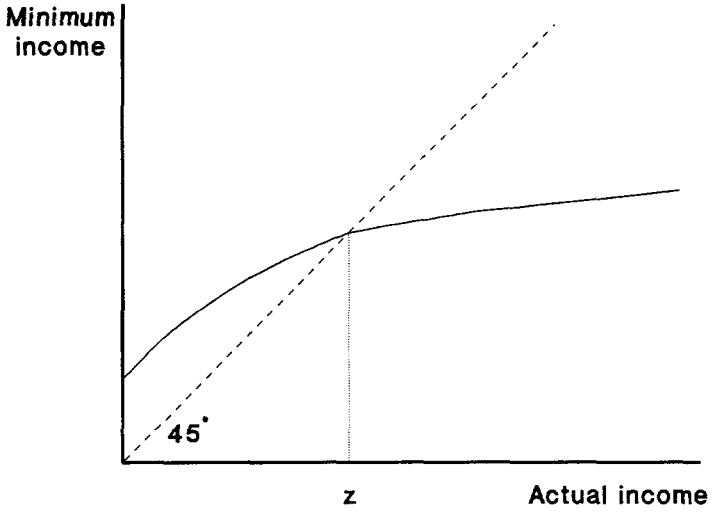

Note: People above z generally deem their own incomes to be adequate, while those below z do not.

approach. Future household surveys for developing countries, such as the LSMS, should consider including subjective poverty line questions.

Dual Poverty Lines

Good practice in any of the approaches to setting a poverty line discussed above is to consider at least two possible poverty lines. The lower one may be interpretable as an ultra-poverty line, such that persons with consumption expenditures below that point are un-able to meet their food energy needs, let alone other food and non-food needs.[27] Indeed, given the arbitrariness involved, there is a good case for considering quite a wide range of the whole distribution of consumption or income. This is the basic idea behind the "dominance approach" which will be discussed in more detail in section 2.6.

27 On the concept of "ultra-poverty" see Lipton (1983, 1988).

By combining the concepts of "absolute" and "relative" poverty lines, there is one straightforward way of setting multiple poverty lines in making poverty comparisons which has a lot of appeal in my view (though it has not yet been used to my knowledge). In each of the situations being compared, say each of two dates, one sets two poverty lines, one of which is fixed in terms of the living standards indicator across both dates, while the other is a relative poverty line reflecting any changes in overall living standards and, hence, perceptions of what constitutes "poverty" in the society. Thus, across any two dates, one can always make distinct comparisons of the changes in both absolute and relative poverty.

2.5 Adding Up Poverty

Let us assume now that a measure of individual well-being has been chosen, and estimated for each person or household in a sample, and that the poverty line is known. How do we aggregate this information into a measure of poverty for each of the distributions being compared?

Poverty Measures

There is now a large literature on poverty measures.[28] Rather than discuss all of the measures that have been used or proposed, I shall focus on three main measures, all of which are members of the class of measures proposed by Foster, Greer and Thorbecke (1984). They are: the head-count index H, the poverty-gap index PG, and the Foster-Greer-Thorbecke P_2 measure. I shall discuss the pros and cons of each. However, rather than treat these as alternative measures, I prefer to interpret them as measures of three different things: the head-count index is a measure of the prevalence of poverty, the poverty-gap index is a measure of the depth of poverty, while the P_2 measures the severity

28 For useful surveys see Foster (1984) and Atkinson (1987).

of poverty. The reasons for these interpretations will soon become clear.

The simplest (and still most common) measure is the <u>head-count index of poverty</u>, given by the proportion of the population for whom consumption (or another suitable measure of living standard) y is less than the poverty line z. Suppose q people are poor by this definition in a population of size n. Then the head-count index is

(2) $H = q/n$ = proportion of total population deemed to be poor.

Is this a good measure of poverty? For some purposes, yes. It is easily understood and communicated. And for certain sorts of poverty comparisons, such as assessing overall progress in reducing poverty, it may be quite adequate (though preferably always calculated for at least two poverty lines, as discussed in the previous section). However, for some purposes, including analyses of the impacts on the poor of specific policies, the head-count index has a serious drawback. To see why, suppose that a poor person suddenly becomes very much poorer. What will happen to measured poverty? Nothing. The head-count index is totally insensitive to differences in the depth of poverty.

A better measure is the <u>poverty gap</u>, based on the aggregate poverty deficit of the poor relative to the poverty line. This gives a good indication of the depth of poverty, in that it depends on the distances of the poor below the poverty line.

To see how this measure is defined, let consumptions be arranged in ascending order, the poorest has y_1, the next poorest y_2, etc., with the least poor having y_q, which is (by definition) no greater than the poverty line z. Then the poverty gap index can be defined as follows:

(3) $$PG = \frac{1}{n} \sum_{i=1}^{q} \left[\frac{z - y_i}{z} \right]$$ = mean proportionate poverty gap across the whole population (zero gap for the non-poor)

This can also be written as:

$$PG = I.H$$

where I is often referred to as the "income gap ratio", and is defined by:

$$(4) \qquad I = \frac{z - y^p}{z} = \text{mean depth of poverty as a proportion of the poverty line}$$

where y^p denotes the mean consumption of the poor. However the income gap ratio defined by (4) is not a good poverty measure. To see why, suppose that someone just below the poverty line is made sufficiently better off to escape poverty. The mean of the remaining poor will fall, and so the income gap ratio will increase. And yet one of the poor has become better off, and none are worse off; one would be loathe to say that there is not less poverty, and yet that is what the income gap ratio would suggest.[29] This problem does not arise if the income gap ratio is multiplied by the head count index to yield PG; under the same circumstances, that measure will register a decrease in poverty.

PG also has an interpretation as an indicator of the potential for eliminating poverty by targeting transfers to the poor. The <u>minimum</u> cost of eliminating poverty using targeted transfers is simply the sum of all the poverty gaps in a population; every poverty gap is filled up to the poverty line. The cost would be

$$\sum_{i=1}^{q} \left[z - y_i \right]$$

Clearly this assumes that the policy-maker has a lot of information; I would not be surprised to find that even a very "pro-poor" government would spend far more than this in the name of poverty reduction. At the

---

29 Note that this does not happen if one calculates the mean consumption of the poorest p% of the population where p is fixed when making the poverty comparison (rather than using the mean below a fixed consumption level).

other extreme, one can consider the maximum cost of eliminating poverty, assuming that the policy maker knows nothing about who is poor and who is not. Then the policy maker would have to give z to everyone to be sure that none are poor; the cost is z.n. By inspecting equation (3) it can be seen that the ratio of the minimum cost of eliminating poverty with perfect targeting to the maximum cost with no targeting is simply PG. Thus this poverty measure is also an indicator of the potential saving to the poverty alleviation budget from targeting.

One drawback of the poverty gap measure is that it may not convincingly capture differences in the severity of poverty. For example, consider two distributions of consumption for four persons; the A distribution is (1,2,3,4) and the B is (2,2,2,4). For a poverty line z=3 (so that H=.75 in both cases), A and B have the same value of PG=.25. However, the poorest person in A has only half the consumption of the poorest in B. The poverty gap will be unaffected by a transfer from a poor person to someone who is less poor.

Sen (1976, 1981) has proposed a better measure of the severity of poverty.[30] However, this measure does not satisfy another useful property which I will simply call "additivity": this requires that aggregate poverty be equal to the population weighted sum of poverty levels in the various sub-groups of society.[31] There are a number of advantages to additivity in the construction of poverty profiles and in testing hypotheses about poverty comparisons, which I will discuss further below.

A measure of the severity of poverty which is additive is the Foster-Greer-Thorbecke P_2 measure, whereby the poverty gaps of the poor

---

30 Also see Kakwani's (1980b) generalization of Sen's measure.

31 The term "additivity" is sometimes also used to describe aggregate measures which are positively weighted sums of sub-group measures but in which those weights are not population shares. I shall not consider such measures here. The term "additive decomposability" is often also used to describe the measures which I have chosen to simply call "additive".

are weighted by those poverty gaps in assessing aggregate poverty. Thus:

$$(5) \qquad P_2 = \frac{1}{n} \sum_{i=1}^{q} \left[\frac{z - y_i}{z} \right]^2 = \text{mean of squared proportionate poverty gaps}$$

In the above example of A and B distributions, the value of P_2 is 0.14 for A and 0.08 for B, indicating the greater severity of poverty in A.

While this measure has clear advantages for some purposes, such as comparing policies which are aiming to reach the poorest, it is not easy to interpret. The measure can be thought of as the sum of two components: an amount due to the poverty gap, and an amount due to inequality amongst the poor. Letting CV_p^2 denote the squared coefficient of variation of consumption amongst the poor, the formulae for P_2 can also be written as the sum of two components:

$$(6) \qquad P_2 = \underbrace{\frac{PG^2}{H}}_{\text{contribution of poverty gap to } P_2} + \underbrace{\frac{(H-PG)^2}{H} \cdot CV_p^2}_{\text{contribution of inequality amongst the poor to } P_2}$$

This interpretation may help, but clearly P_2 is not as easy to interpret as PG or (especially) H. That is a drawback for expository purposes. For poverty comparisons, however, the key point is that a ranking of dates, places, or policies in terms of P_2 should reflect well their ranking in terms of the severity of poverty. It is the ability of the measure to order distributions in a better way than the alternatives that makes it useful, not the precise numbers obtained.

On comparing the above formulae for H, PG and P_2 a common structure is evident. This suggests a generic class of additive measures:

$$(7) \qquad P_\alpha = \frac{1}{n} \sum_{i=1}^{q} \left[\frac{z - y_i}{z} \right]^\alpha$$

for some non-negative parameter α. This is the Foster-Greer-Thorbecke class of poverty measures (Foster et al., 1984). P_α is simply the mean over the whole population of an individual poverty measure which takes the value $(1 - y_i/z)^\alpha$ for the poor and zero for the non-poor. (Which is also the simplest and most accurate way to calculate P_α if one has access to the individual or household level data; more common forms of grouped data are discussed below.) The main measures discussed so far are special cases; the head-count index has α=0, while α=1 for PG. For both the poverty-gap index and P_2 the individual poverty measure is strictly decreasing in the living standard of the poor (the lower the standard of living the poorer you are deemed to be). Furthermore, P_2 has the property that the increase in your measured poverty due to a fall in standard of living will be deemed greater the poorer you are.

Figure 3 shows how the relationship between individual poverty and standard of living varies across the different values of α.[32] The higher the value of α the more sensitive the measure is to the well-being of the poorest person; as α approaches infinity the measure collapses to one which only reflects the poverty of the poorest person.

---

32 This also illustrates another conceptual attraction of P_2, namely that the weight hits zero smoothly at the poverty line; thus there is negligible difference to the weight the measure attaches to someone just above the poverty line versus someone just below it. Given the uncertainties in measuring living standards discussed above, this is a desirable property.

Figure 3: Individual Poverty Measures

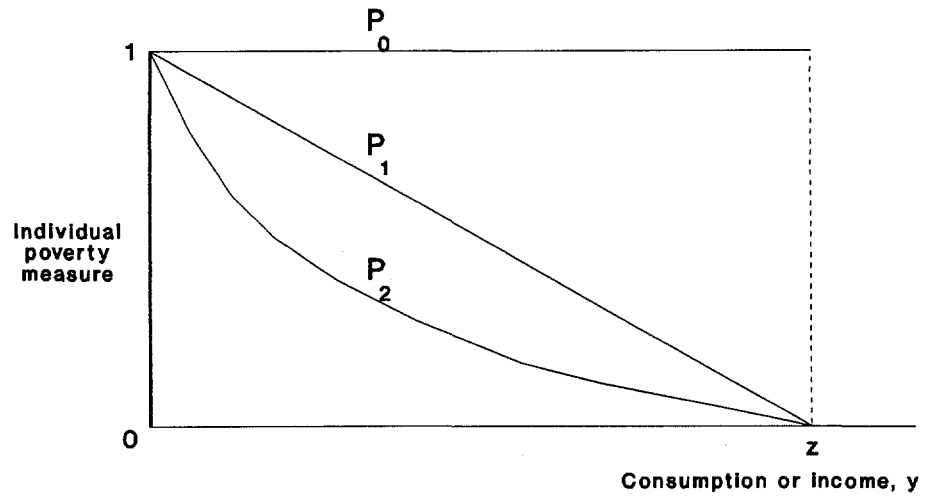

Note: The figure gives the individual poverty measures implied by various P_α poverty meaures.

There are other additive poverty measures. One can characterize a general class of additive measures taking the following form:

$$P = \sum_{i=1}^{n} p(z, y_i)/n \tag{8}$$

where $p(z, y_i)$ 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. On the specific functional forms assumed by other measures in the general class defined by equation (8) see Atkinson (1987).

Given the intellectual energy that has gone into the theory of poverty measurement over the last 15 years, and the virtual plethora of poverty measures from which one can now choose, it is of interest to ask: Does it really matter in poverty comparisons which of these measures one uses?

The answer depends on whether, and how, relative inequalities in the society have changed across the situations being compared. If all consumption levels (poor and non-poor) have changed by the same proportion - a "distributionally neutral" growth or contraction - then all of these poverty measures will yield the same ranking in the poverty comparison, and the ranking in terms of absolute poverty will depend solely on the direction of change in the mean of the distribution. This is evident if one returns to equation (1) in which the parameters **L** capture relative inequalities (the Lorenz curve). Rankings will also be unaffected for sufficiently small changes in relative inequalities.

However, the differences between these measures can become quite pronounced otherwise. 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; an example is given in section 3.7.) Policy B entails the opposite change - the poorest lose while those at 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; $H_A > H_B$ since changes in H depend solely on which direction people are crossing the poverty line. However, a measure such as P_2 will indicate the opposite ranking, $P_{2A} < P_{2B}$, since it will respond relatively more to the gains amongst the poorest than amongst the not-so-poor.

It is often important to know whether a poverty comparison is sensitive to the choice of poverty measure, and not just because of the uncertainties involved in that choice - differences in rankings by different measures are also telling us something about the precise way in which the distribution of living standards has changed. Section 2.6

will return to this point, and outline some analytical tools which can help assess sensitivity to the choice of poverty measure.

Estimation

It is sometimes said that one rarely has access to the data needed to estimate the more "sophisticated" poverty measures discussed above, and so the best one can typically do is estimate the head-count index. This is not true. I have never seen a data set which would only permit estimation of the head-count index; indeed, I have never seen one for which the marginal cost of estimating the poverty gap (or other FGT measures) over the head-count index is anything but negligible with even modest computing resources.

Two distinct types of data set are encountered in practice: individual ("unit") record data and tabulated grouped data, of which the latter is the more common. The former is typically only available in a machine readable form (typically magnetic tape), while the latter is often found in governmental statistical publications. Quite different problems are encountered in estimating poverty measures on these two types of data.

All of the additive poverty measures above can be readily and accurately calculated as the means of the corresponding individual poverty measures when one has access to unit record data. The main points to be aware of are:

i) It should not be presumed that estimates from unit record data are more accurate than those from grouped data, since the latter can "average out" errors in the unit-record data, such as negative consumption figures (which might otherwise add a sizable bias to estimates of the severity of poverty).

ii) Most large household surveys use stratified samples, whereby the chance of being selected in the sample is not uniform over the population. This is often done to assure adequate sample sizes in

certain regions (see section 2.2). Estimates of population parameters from a stratified sample are only un-biased if weighted by the appropriate inverse sampling rates. Provided your data set includes the sampling rate for each household or area, it is easy to correct for this potential source of bias; for further discussion see, for example, Levy and Lemeshow (1991).

iii) One should be clear on whether one wants to estimate poverty amongst households or poverty amongst people. For example, suppose one ranked households by consumption per person, but measured poverty in terms of the proportion of households who are below the poverty line. Household size tends to be negatively correlated with consumption per person, so your calculation will tend to underestimate the number of persons who are living in poor households (though it will not necessarily underestimate the number of poor persons; that also depends on the distribution within the household, which is typically unknown).

In my view, the most defensible position on this last issue is to recognize that poverty is experienced by individuals not households per se, and so it is poverty amongst persons that we are trying to measure. Though we may not know anything about distribution within the household, that does not mean that we should only measure poverty amongst households. A common practice is to assume a uniform distribution within households when constructing the estimated distribution of individual consumptions. This may well lead one to underestimate poverty amongst persons, and the magnitude need not be negligible (Haddad and Kanbur 1990). However, it is not clear what would be a better assumption.

While individual or household level data are the ideal, in practice we often have to make do with tabulations of grouped data, such as income shares of households ranked in deciles, or income frequency distributions. (And even household level data can be interpreted as grouped individual data.) Poverty lines rarely occur at the boundaries

of the grouped data. And so we must find some way of interpolating between those boundaries.

Linear interpolation is the easiest and most common method, but it can be quite inaccurate, particularly when the poverty line is far from the mode of the distribution (such as when it is found in the typically quite non-linear lower region of the cumulative frequency distribution.[33]) Quadratic interpolation is usually feasible with the same data, and is generally more accurate, though one must be wary of the possibility that the probability density (slope of the frequency distribution) implied by this method does not become negative.

A method of interpolation which can be very accurate, and is also useful in certain policy simulations, involves estimation of a parameterized Lorenz curve. There are many specifications to choose from. Though it is easy to obtain a good fit for a smooth function such as a Lorenz curve, that does not mean that it will give accurate estimates of poverty measures. The accuracy depends greatly on the particular specification used, and some tend to dominate others on many data sets. The two best in my experience are the generalized quadratic model (Villasenor and Arnold 1989) and the Beta model (Kakwani 1980b). Formulae for the above poverty measures as functions of the Lorenz curve parameters for two specifications and the mean of the distribution have been derived by Datt (1991).[34]

Table 2 gives illustrative calculations of the head-count index of poverty for 1984 in Indonesia by various methods using:

---

33 For example, I have come across one estimate of the head-count index for a country (which can remain un-named) which was obtained by linear interpolation in the lowest class interval of the distribution; the estimate was 9.5%. However, when one re-estimated with a model (based on a Beta specification; see below) which took account of the non-linearity one obtained a figure of 0.5%! This is probably an extreme case, though large errors are likely from linear interpolation at the lower end of any grouped distribution.

34 The program POVCAL (Chen, Datt and Ravallion, 1991) performs these calculations. It is available as a self-contained user friendly program for use on any IBM compatible PC.

i) the primary data tapes of Indonesia's socio-economic survey (SUSENAS) of consumption for 1984 covering 55,000 sampled households (weighted by the appropriate inverse sampling rates),

ii) parameterized Lorenz curves (using the Beta model) calibrated to a detailed description of the frequency distribution based on 50 class intervals formed from the unit record data, of which 18 intervals are below the poverty line,

iii) Lorenz curves calibrated to far more "coarse" frequency distributions using 15, 10, and 5 class intervals, of which 8, 4, and 2 respectively are below the poverty line. These are typical of the detail usually available from published sources. With the more aggregated data it is also very common for the poverty line to fall well within a class interval; I have constructed grouped distributions in which the poverty line is in the middle of a class interval.

Table 2: Alternative Methods of Estimating Poverty Measures for Indonesia

Poverty measure	Calculated directly from unit record data (1984, sample=55,000)	Number of intervals used in estimation from grouped data: 50 (18 "poor")	15 (8 "poor")	10 (4 "poor")	5 (2 "poor")
Head-count index H (%)	33.02	33.74	33.64	33.88	33.63
Poverty gap index PG (%)	8.52	9.10	9.04	9.17	9.10

Source: Ravallion, Datt and van de Walle (1991).

The results in Table 2 confirm a loss of accuracy in using grouped data, though it is not large; the head-count index, for example, is estimated within three-quarters of a percentage point using the grouped data, though the proportionate error is higher for the poverty gap index. Furthermore, the loss of accuracy is affected much by

contraction in the number of class intervals in the grouped data. I do not know how specific this result is to these data. However, it does suggest that there need not be much loss of accuracy in using even highly aggregated data for measuring poverty.

Poverty measures can be quite sensitive to certain sorts of measurement error in the underlying parameters, and quite robust to others. Consider, for example, errors in the mean of the distribution on which the poverty measure is being estimated. It is easily demonstrated that the elasticity of the head-count index of poverty to errors in the mean holding the Lorenz curve constant is simply the elasticity of the cumulative distribution function evaluated at the poverty line. This is also easily estimated, and (from my experience) values around two are quite common for developing countries.[35] (Furthermore, amongst the FGT class of measures, the elasticities tend to be higher the higher the value of α.[36]) 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.

Note that this assumes that there is no corresponding error in the Lorenz curve. It is generally thought to be more likely that survey measurement errors will tend to lead to an under-estimation of the extent of inequality; this will be so if the rich tend to underestimate proportionately more than the poor. Indeed, it may well be that surveys overestimate consumption of the poor (through social stigma facing respondents or because the sample frame excludes sub-groups of the poor,

---

35 For example, Ravallion and Huppi (1991) estimated that the elasticity of the rural head-count index with respect to the mean was -2.0 in 1984, while the urban figure was -3.3.

36 For example, while the national head-count index for Indonesia in 1984 was estimated to have an elasticity with respect to the mean of -2.1, the PG measure had an elasticity of -2.9, and P_2 had an elasticity of -3.4 (Ravallion and Huppi 1991). The same pattern has been observed in my estimates for other developing countries.

such as the homeless, who are amongst the poorest), and underestimate that of the rich (such as through their fear of revealing black market transactions or tax evasion). But, it is less clear how this will translate into poverty measurement errors. For example, if the underestimation of income or consumption was entirely above the poverty line then it would make no difference to measured poverty. More than this, there is little that one can say in general about the effects on the head-count index of measurement error in the Lorenz curve when the estimated consumptions of the poor are affected.[37] However, for the FGT measures PG and P_2 survey underestimation (overestimation) of consumption by the poor will lead to overestimation (underestimation) of poverty. Furthermore, the higher the value of α the more sensitive the measure will be to consumption measurement errors for the poorest.

In making quantitative poverty comparisons over time, the rate of inflation is one source of error in estimates of mean real consumption per capita which does affect all consumption levels by the same proportion (and hence leaves the Lorenz curve unaffected). However, this error affects both the mean and the poverty line, and so (recalling that most poverty measures are homogeneous in the poverty line and mean as in equation 1) measures of poverty will be unaffected.

Hypothesis Testing

Testing hypotheses about differences in poverty between two situations is not difficult for additive poverty measures. Recall that additive measures can be calculated as the sample mean of an appropriately defined individual poverty measure. For random samples,

---

37 For example, when the poverty line is at the mean of the distribution, the head-count index will be unaffected by errors at the same percentage along all points along the Lorenz curve.

the standard error can then also be readily calculated.[38] This allows us to test hypotheses about poverty, such as whether it is significantly higher in one sub-group than another.

For the head-count index, the standard error can be calculated the same way as for any population proportion.[39] Thus the standard deviation of the sample distribution of the head-count index is $\sqrt{\{H.(1-H)/n\}}$ for a sample of size n, and so, by the properties of the normal distribution, there is (for example) a 95% chance that the true value of the head-count index lies in the interval:

$$(9) \quad H - 1.96\sqrt{\{H.(1-H)/n\}} < H < H + 1.96\sqrt{\{H.(1-H)/n\}}$$

To test the null hypothesis that $H_A = H_B$ in distributions A and B with sample sizes n_A and n_B one should calculate the test statistic $t = (H_A - H_B)/s$ where s denotes the standard deviation of the sampling distribution of $H_A - H_B$ which (under the null hypothesis) is given by:

$$(10) \qquad s = \sqrt{\{H.(1-H).(\frac{1}{n_A} + \frac{1}{n_B})\}}$$

where $H = (n_A H_A + n_B H_B)/(n_A + n_B)$; see, for example, Hamburg (1977). If the calculated value of t has an absolute value less than 1.96 (2.58) then the difference in head-count indices between two dates, say, cannot be considered statistically significant at the 5% (1%) level using a two tailed test.

---

38 The key result from statistics is the Central Limit Theorem. Let a sample mean m of any variable be calculated from a random sample of size n, and let μ denote the true value of the mean. Then $(m-\mu)\sqrt{n}$ approaches a normal distribution as n increases.

39 Which will have a binomial distribution in random samples, approaching a normal distribution as sample size increases. For all but very small sample sizes (n<5), a useful rule-of-thumb is that the approximation involved in using the normal distribution will be accurate as long as the absolute value of $\sqrt{\{(1-H)/H\}} - \sqrt{\{H/(1-H)\}}$ does not exceed $0.3\sqrt{n}$ (Box et al., 1978).

These methods can be extended to other additive poverty measures. Kakwani (1990) has derived useful formulae for the standard errors of a number of other additive measures, including the Foster-Greer-Thorbecke P_α measures. (Also see Blundell and Preston, 1991.) The standard error of the P_α measure is $\sqrt{\{(P_{2\alpha}-P_\alpha^2)/n\}}$ which yields the aforementioned standard error for the head-count index as the special case when $\alpha=0$.

2.6 Decompositions

Decompositions can be useful analytical tools in making quantitative poverty comparisons. I will first discuss how a single aggregate poverty number can be decomposed to form a poverty profile. I shall then look at two ways of decomposing changes in poverty over time.

Poverty Profiles

A "poverty profile" is simply a special case of a poverty comparison, showing how poverty varies across sub-groups of society, such as region of residence or sector of employment. A poverty profile can be extremely useful in assessing how the sectoral or regional pattern of economic change is likely to affect aggregate poverty.[40] For example, if the poverty profile shows that there is significantly more poverty in the rural farm sector than the non-farm sector then a policy reform which improves farmers' terms of trade is very likely to reduce aggregate poverty in the economy as a whole (even if it leads to higher poverty in the non-farm sector).

Additive poverty measures, such as the FGT class, can greatly facilitate poverty comparisons. Consider the general class of additive poverty measures described by equation (8). Suppose the population can be divided into m mutually exclusive sub-groups forming the poverty profile. The poverty profile is simply the list of poverty measures P_j

---

40 See Kanbur (1987, 1990) for a discussion of the uses of poverty profiles in policy analysis.

for j=1,..,m. Aggregate poverty can then be written as the population weighted mean of the sub-group poverty measures:

$$(11) \qquad P = \sum_{j=1}^{m} P_j n_j / n$$

where

$$P_j = \sum_{i=1}^{n_j} p(z_j, y_{ij})/n_j$$

is the poverty measure for the j'th sub-group with population n_j, having consumptions y_{ij} for $i=1,..,n_j$, and the total population is $n=\Sigma n_j$. 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 of additive poverty measures in forming poverty profiles, additivity guarantees that when poverty increases (decreases) in any sub-group of the population, aggregate poverty will also increase (decrease) (Foster et al., 1984; Foster and Shorrocks 1991). The latter property is intuitively sensible for any poverty profile. Indeed, an evaluation of the effects on aggregate poverty of targeted poverty alleviation schemes - whereby the benefits are concentrated in certain sub-groups - may be quite misleading unless the poverty measure used has this property; the measure of aggregate poverty may show an increase even if poverty fell in the target group, and no others lost. There are otherwise attractive measures in the literature which can fail to satisfy this condition, such as the Sen and Kakwani indices.[41]

One of the main problems encountered in forming a reliable poverty profile has already been discussed, namely that of setting the poverty lines for different sub-groups. Section 3.3 will discuss further the

---

41 See Foster and Shorrocks (1991) who characterize the general class of sub-group consistent poverty measures.

problems encountered in constructing reliable poverty profiles using empirical examples.

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. 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".[42] The type A profile is not always a feasible option, notably in situations where the population cannot be classified into mutually exclusive groups. For example, data on consumption of various commodities by different income groups cannot be presented in such a form, though it can be presented as a type B poverty profile. But in many circumstances the poverty profile can be presented in either form; which is preferable?

Consider the hypothetical data in the left hand panel of Table 3. The economy consists of 1000 people living in two regions, "north" and "south". From a household sample survey we can estimate the numbers of "poor" and "non-poor" in each region, as given in Table 3. The two types of poverty profile described above are then given in the right hand panel of the Table. They clearly give a very different impression, though of course they are measuring quite different things.

---

42 Both types are common in poverty studies; for examples and further discussion see van de Walle (1991).

Table 3: Alternative Representations of a Poverty Profile for Hypothetical Data

Region	Numbers of persons		Poverty profile	
	Poor	Non-poor	**Type A** Percent of region's population who are poor	**Type B** Percent of total poor in each region
South	100	100	50	33
North	200	600	25	67

The type A poverty profile is likely to be more useful for policy. Suppose that (to give a stylized version of a common policy problem), 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. This is an example of what is sometimes called "indicator targeting". It is "imperfect" targeting because (as is invariably the case) the policy maker does not know who has which standard of living even when a distribution of living standards can be constructed from a household sample survey; rather the policy maker relies on an imperfect indicator of living standards, in this case region of residence; section 3.8 will discuss this class of policy problems in further detail.

A moment's reflection will confirm that more of that money will go to the poor if the "south" is chosen as the target region for the data in Table 3. The type A profile is a better guide for targeting when one is aiming to have greatest impact on the poverty gap. This is an example of a quite general principle: when making lump-sum transfers to different sub-groups of a population with the aim of minimizing the aggregate value of the FGT poverty measure P_α the next unit of money should go to the sub-group with the highest value of $P_{\alpha-1}$ (Kanbur 1987).

Another way to represent the information in a poverty profile is to normalize each sub-group's poverty measure by the aggregate ("national") poverty measure.[43] This will make no difference to the qualitative poverty comparison for a given date (say), though it could be misleading when comparing poverty profiles over time if the level of aggregate poverty has also changed over time. Section 3.6 will discuss a way of assessing the contribution of different sub-groups in the poverty profile to changes over time in aggregate poverty.

Decomposing a Change in Poverty: Growth and Redistribution Components

It is sometimes of interest to ask: how much of any observed change in poverty can be attributed to changes in the distribution of living standards, as distinct from growth in average living standards? The usual inequality measures, such as the Gini index, can be misleading in this context. One certainly cannot conclude that a reduction in inequality (by any measure satisfying the transfer principle mentioned in section 2.3) will reduce poverty. And even when a specific reduction (increase) in inequality does imply a reduction (increase) in poverty, the change in the inequality measure can be a poor guide to the quantitative impact on poverty. A time series of an inequality measure can be quite uninformative about how changes in distribution have affected the poor.

Datt and Ravallion (1991b) propose a simple decomposition for any change in measured poverty which allows one to rigorously quantify the relative importance of growth versus re-distribution. The change in poverty is decomposed as the sum of a growth component (the change in poverty that would have been observed if the Lorenz curve had not shifted), a redistribution component (the change that would have been

---

43 For additive poverty measures, the normalized poverty measure for sub-group j, P_j/P, can be interpreted as a measure of "poverty intensity"; see Rodgers and Rodgers (1991).

observed if the mean had not shifted), and a residual (the interaction between growth and redistribution effects).

Recalling the notation used in equation (1), let $P(z/\mu,L)$ denote measured poverty when the distribution of living standards has the mean μ and Lorenz curve L. The change in poverty between dates 1 and 2 (say) can then be decomposed as follows:

$$(12) \qquad P_2 - P_1 = \underbrace{G(1,2;r)}_{\text{growth component}} + \underbrace{D(1,2;r)}_{\text{redistribution component}} + \underbrace{R(1,2;r)}_{\text{residual}}$$

in which the growth and redistribution components are defined by

$$G(1,2;r) \equiv P(z/\mu_2,L_r) - P(z/\mu_1,L_r)$$
$$D(1,2;r) \equiv P(z/\mu_r,L_2) - P(z/\mu_r,L_1)$$

while R( ) in (12) denotes the residual. In each case, the first two arguments in the parentheses refer to the initial and terminal dates of the decomposition period, and the last argument makes explicit the reference date r with respect to which the observed change in poverty is decomposed. The residual in (12) exists whenever the poverty measure is not additively separable between μ and L, i.e., whenever the marginal effects on the poverty measure of changes in the mean (Lorenz curve) depend on the Lorenz curve (mean). In general, the residual does not vanish. It can be interpreted as the difference between the growth (redistribution) components evaluated at the terminal and initial Lorenz curves (mean incomes) respectively. If the mean income <u>or</u> the Lorenz curve remains unchanged over the decomposition period, then the residual vanishes.

The Sectoral Decomposition of a Change in Poverty

When analyzing the sources of observed reductions in aggregate poverty, we can also make use of another simple decomposition formula, proposed in Ravallion and Huppi (1991), and also exploiting the additivity property of the FGT class of measures, as discussed in section 2.5. The idea here is to throw light on the relative importance of changes within sectors versus changes between them, such as due to inter-sectoral population or work-force shifts.

To see how this can be done, let P_{it} denote the FGT poverty measure (or any other additive measure) for sector i with population share n_i at date t, where there are m such sectors, and t=1, 2. Then it is readily verified that:

$$(13) \quad P_2 - P_1 = \Sigma(P_{i2} - P_{i1})n_{i1} \quad \text{(Intra-sectoral effects)}$$

$$+ \Sigma(n_{i2} - n_{i1})P_{i1} \quad \text{(Population shift effects)}$$

$$+ \Sigma(P_{i2} - P_{i1})(n_{i2} - n_{i1}) \quad \text{(Interaction effects)}$$

where all summations are over i=1,..,m. The "intra-sectoral effects" tell us the contribution of poverty changes within sectors, controlling for their base period population shares, while the "population shift effects" tell us how much of the poverty in the first date was reduced by the various changes in population shares of sectors between then and the second date. The interaction effects arise from the possible correlation between sectoral gains and population shifts; the sign of the interaction effect tells us whether people tended to switch to the sectors where poverty was falling or not.

2.7 The Robustness of Qualitative Poverty Comparisons

At a number of points in the discussion so far we have seen that there is pervasive uncertainty about possibly crucial aspects of a poverty comparison. There are likely to be errors in our living

standards data, unknown differences in needs between households at similar consumption levels, uncertainty and arbitrariness about both the poverty line and precise poverty measure. Given these problems, how robust are our poverty comparisons? Would they alter if we made alternative assumptions?

A recent strand of research in poverty analysis has shown how we can answer such questions, drawing on and developing results from the theory of stochastic dominance. I shall give an elementary exposition of the approach, again oriented to the needs of the analyst trying to make a reasonably robust poverty comparison.[44] The analysis is easier for a single dimension of well-being, but I will also give an introduction to multi-dimensional dominance.

A Single Measure of Standard of Living

Consider again the three poverty measures discussed in section 2.5, namely H, PG and P_2. Imagine the curve which is traced out as one plots H 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, which can be thought of as the "poverty incidence curve" - each point on the curve gives the proportion of the population consuming less than the amount given on the horizontal axis (Figure 4, panel (a)). If one calculates the area under this curve up to each point then one traces out the "poverty deficit curve" (Figure 4, panel (b)). Each point on this curve is simply the value of PG times the poverty line z. If one again calculates the area under the poverty deficit curve at each point then one obtains a new curve, which can be termed the "poverty severity

44 On the theory of stochastic dominance see Quirk and Saposnik (1962), Hadar and Russell (1969), and Rothschild and Stiglitz (1970). On the use of dominance conditions in ranking distributions in terms of measures of inequality see Atkinson (1970); on rankings in terms of poverty see Atkinson (1987), and Foster and Shorrocks (1988).

curve" (Figure 4, panel (c)); each point on this curve is directly proportional to P_2. The Appendix gives more formal definitions.

Suppose we do not know the poverty line z, but we can be sure that it does not exceed z^{max}. Nor do we know the poverty measure, but we can identify some desirable properties for such a measure, including the aforementioned additivity property.[45] Then it can be shown that poverty will unambiguously fall between two dates if the poverty incidence curve (the cumulative distribution) for the latter date lies nowhere above that for the former date, up to z^{max}. This is called the First-Order Dominance Condition (FOD).

Figure 4, panel (a), illustrates first-order dominance. When we plot the cumulative frequency distributions (cumulative percentages of the population below various consumption levels) in states A and B, we find that the curve for A is everywhere above that for B. Poverty is higher in state A than in B, no matter what the poverty line or measure.

If the curves intersect as in panel (a) of Figure 5 (and they may intersect more than once), then the ranking is ambiguous. Then we know that some poverty lines and some poverty measures will rank the distributions differently to others. We need more information. One can restrict the range of poverty lines, or one can impose more structure on the poverty measure.

If one restricts attention to additive measures which do reflect the depth of poverty such as PG and P_2 (but excluding the head-count index) then we can use a Second-Order Dominance Condition (SOD). A fall in poverty then requires that the poverty deficit curve, given by the

---

45 More precisely, attention is restricted to poverty measures which are additive, of the form in (8), or any measure which can be written as a monotonic transformation of an additive measure. All the FGT measures discussed in section 2.5 qualify. Atkinson (1987, 1989, Chapter 2) characterizes the set of admissible poverty measures and gives other examples from the literature.

Figure 4: The Construction of the Three Poverty Curves

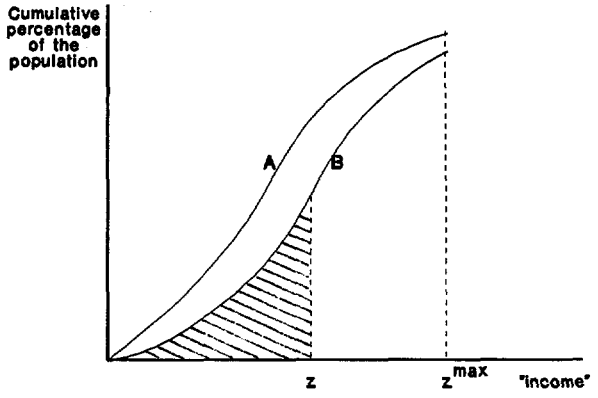

(a) Poverty Incidence Curves for two Distributions A and B

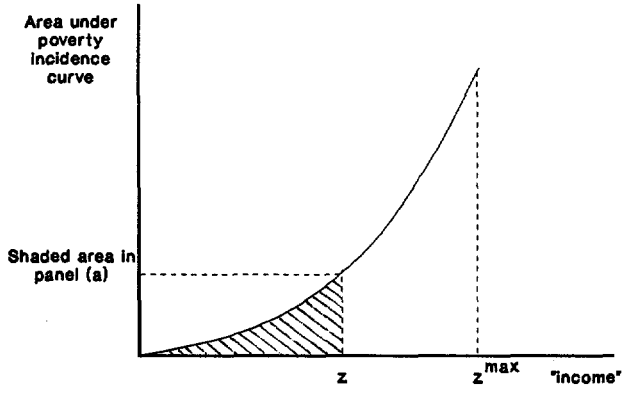

(b) Poverty Deficit Curve for Distribution B

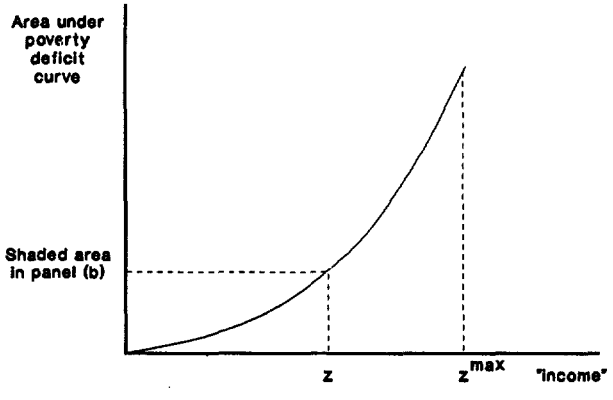

(c) Poverty Severity Curve for Distribution B

Figure 5: Intersecting Poverty Incidence Curves

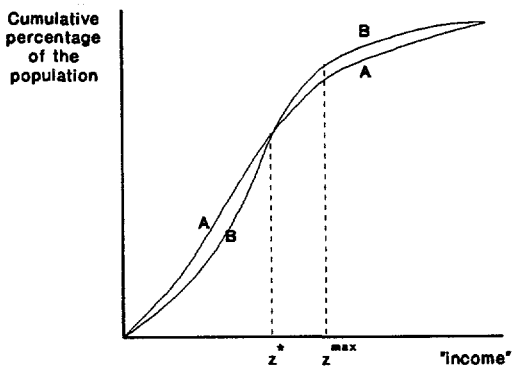

(a)

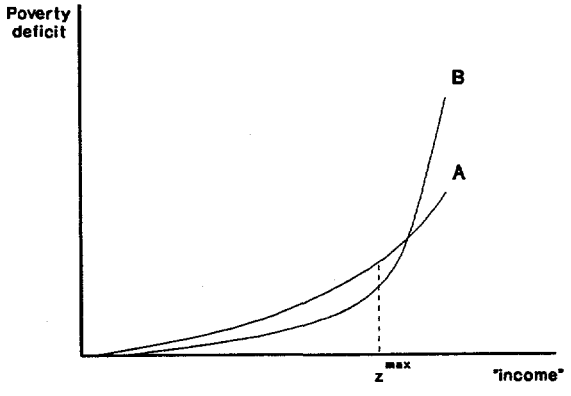

(b)

Note: Poverty is higher for the A distribution in panel (a) if the maximum admissible poverty line is less than z^*; this holds for all three poverty measures, H, PG and P_2 (and others). Suppose the poverty line could be as high as z^{max}. Poverty is then higher for A if its poverty deficit curve (the area under the curve in panel (a)) is above that for B up to z^{max}, as in panel (b); this holds for PG and P_2, but not H.

area under the cumulative distribution, is nowhere lower for the earlier date at all points up to the maximum poverty line, and at least somewhere higher.[46] This is illustrated in panel (b) of Figure 5.

When this test is inconclusive, one can further restrict the range of admissible poverty measures. If one is content to rely solely on distribution sensitive measures such as P_2 (but now excluding H and PG) then a Third-Order Dominance Condition (TOD) can be tested; an unambiguous poverty comparison for all poverty lines then requires that the poverty severity curve is everywhere higher in one of the two situations being compared. If necessary, one can go on to test higher order dominance, though the interpretation of the (increasingly) restricted class of measures becomes less clear.[47]

To illustrate the three dominance tests, consider an initial state in which three persons have consumptions in amounts (1, 2, 3). Any final state in which one or more of these persons has a higher consumption, and none have a lower consumption, will imply a lower poverty incidence curve (strictly lower and no higher anywhere), and hence no higher poverty for any poverty line or poverty measure; examples of such final states are (2, 2, 3) or (1, 2, 4). Consider instead the final state (2, 2, 2). The poverty incidence curves now cross each other: some poverty lines and some poverty measures will judge this final state to be an improvement, while others will judge it to be worse than the initial state. However, the poverty deficit curves

---

46 This is equivalent to generalized Lorenz curve dominance; if the generalized Lorenz curve (ordinary Lorenz curve scaled up by the mean) of distribution A is everywhere above that of B then the area under A's cumulative frequency distribution must be everywhere lower than B's. On the poverty deficit curve see Atkinson (1987). On the generalized Lorenz curve see Shorrocks (1983). On the relationship between the two see Atkinson and Bourguignon (1989).

47 Amongst the FGT class, the forth-order dominance test restricts attention to P_α measures for values of α=3 or higher. See Kakwani (1980b) for an interpretation of such measures.

do not cross each other; for the initial state the poverty deficit curve is (1, 3, 6) (corresponding to consumptions 1, 2, and 3), while it is (0, 3, 6) for the final state. Thus poverty will have fallen (or at least not increased) for all poverty lines and all measures which are decreasing in consumptions of the poor, such as PG and P_2.[48]

But what if the final state has consumptions (1.5, 1.5, 2)? Table 4 gives the poverty incidence, deficit and severity curves. Even if we confine attention to distribution sensitive poverty measures, some poverty lines will rank the states differently to others. But note that the intersection point of the poverty severity curves is above 2; any poverty line less than this point will indicate that poverty has fallen for all distribution sensitive measures, such as P_2.[49]

Table 4: Poverty Incidence, Depth and Severity Curves for Three People with Initial Consumptions (1,2,3) and Final Consumptions (1.5,1.5,2)

Consumption (z)	Poverty Incidence Curve F(z)		Poverty Depth Curve D(z)		Poverty Severity Curve S(z)	
	Initial	Final	Initial	Final	Initial	Final
1	1/3	0	1/3	0	1/3	0
1.5	1/3	2/3	2/3	2/3	1	2/3
2	2/3	1	4/3	5/3	7/3	7/3
3	1	1	7/3	8/3	14/3	15/3

More than One Dimension

Similar ideas can be applied in circumstances when poverty lines vary across households or individuals in an unknown way. For example, errors in measuring the standard of living can entail that we should be

48 There is an unambiguous fall in poverty for all such measures as long as the poverty line is 2 or less.

49 Fourth-order dominance holds for all points, implying that FGT measures for α=3 or higher will show a fall in poverty for all possible poverty lines.

using different 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, FOD of one distribution over another implies an unambiguous ranking in terms of the head-count index of poverty. This holds no matter what the underlying distribution of poverty lines. For further discussion in the context of measuring undernutrition when nutrient requirements vary in some unknown way see Kakwani (1989) and Ravallion (1992a).

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.[50] In the special case in which the

---

50 For a general discussion of multi-variate dominance tests for various assumptions about the multiple dimensions interact in determining welfare see Atkinson and Bourguignon (1982). In the specific context of inequality comparisons when needs differ see Atkinson and Bourguignon (1987) and Bourguignon (1989). The discussion in Atkinson (1988) is in the context of poverty measures.

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 amongst (say) rural households, separately to urban households, or large families separately to 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 SOD for each "needs" group separately, or TOD if necessary.

These will often be quite stringent tests. Weaker tests can be invoked if one is willing to rank needs groups in terms of the marginal welfare attached to an increment of consumption. Suppose one can, and let group 1 have the highest marginal social valuation of consumption (i.e., the steepest individual poverty measure). Following Atkinson and Bourguignon (1987), let us also assume that this ranking is the same at all possible consumption levels (so group 1 always has the highest marginal valuation of consumption). When ranking distributions in terms of poverty measures we will also need to assume that the poverty measure as a function of consumption is not discontinuous at the poverty line (Atkinson 1988). This precludes the head-count index, but few other measures; the condition holds for PG and P_2 as illustrated in Figure 3. Under these conditions one can apply simple <u>partial dominance tests</u>, where the test is done cumulatively by the ranked needs groups starting from group 1, rather than separately for each group (Atkinson and Bourguignon 1987, Atkinson 1988). Thus dominance is tested on the

cumulative frequency distributions for group 1 in two situations compared, then for the population weighted sum of groups 1 and 2, then for 1,2, and 3, etc. This makes dominance more likely. For example, although poverty may increase in some needs groups, aggregate poverty may be found to have fallen as a result of some policy change.

However, these tests have to be modified further when the distribution of needs also changes, such as when the proportion of the population living in urban areas has increased over the period of the poverty comparison, as is typically the case in inter-temporal poverty comparisons for developing countries. It is theoretically possible that FOD may hold separately for each of urban and rural areas, and yet not hold in the aggregate for all possible distributions of needs between the two sectors and all possible ways in which consumption and needs interact to determine well-being. More general tests can be devised for such situations, though they are difficult to explain non-mathematically; see Atkinson and Bourguignon (1982).

When two frequency distributions are quite close, we may also want to assess whether the difference between them is statistically significant. This can be done quite easily using the Kolmogorov-Smirnov test, based on the largest vertical distance between the two cumulative frequency curves; expositions on this simple test, and tabulations of critical values are readily available (see, for example, Daniel 1990 Chapter 8). More advanced methods are needed for testing second or higher-order dominance; see Chow et al (1991). Results are not yet available for hypothesis testing of multi-dimensional dominance, though that may soon change.

In summary: Dominance testing can be a powerful tool for assessing whether poverty has increased over time, or as a result of some policy reform. The tests can be robust to many of the measurement problems that routinely confound such poverty assessments. And the tests are often easy to apply; for the first-order test (which assumes very little

about how poverty is measured, but is thus the hardest test to pass) it is simply a matter of plotting the cumulative frequencies of consumption in each of the situations being compared. When the test is inconclusive, a second-order test may prove useful, based on the area under the plotted cumulative distribution. This restricts attention to measures which reflect the depth of poverty. There are also tests which can handle multiple dimensions of well-being which cannot be precisely aggregated. On-going research on this topic will probably further enhance the undeniable advantages of this approach in situations where data are imperfect and measurement is inherently controversial.

3. Putting Theory into Practice

There are two distinct tasks for which poverty comparisons are called for. The first is as part of an overall assessment of a country's progress in economic development; the second is in the evaluation of specific policies or projects, actual or contemplated. The first task can also be viewed as policy assessment in a broad sense; indeed, given the difficulties often encountered in evaluating specific policies, the best we might reasonably hope for in many situations is a reliable overall assessment of development progress.

This section addresses a number of questions that often arise in poverty assessment and policy work, and tries to illustrate how the concepts and methods discussed in section 2 can throw light on the answers. One cannot hope to cover every issue that might arise in practice. I start with a few of the most common questions about poverty data, and then move on to a few case-studies of questions about the impacts of policies on poverty.[51]

3.1 How Well Can the Prevalence of Poverty in a Country be Predicted Without a Household Survey?

Suppose that we do not have access to a household level budget or income survey, such as discussed in section 2.2. Can we still make a reasonably reliable estimate of the prevalence of poverty in a country on the basis of the more readily available aggregate economic and social indicators, such as available in the World Development Report?

The only way to convincingly answer this question is to first estimate the poverty measure from a household survey, and then try to

---

51 The examples reported here draw solely on my own work (often in collaboration with others); most of this was done as part of a World Bank research project, "Policy Analysis and Poverty: Applicable Methods and Case Studies". Thus it is by no means a comprehensive survey of recent analytical work on poverty in developing countries, and the coverage is decidedly biased toward Asia. However, the aim is not to present such a survey, but rather to illustrate the methods outlined in section 2.

predict its value from the readily available aggregate indicators, pretending that we do not have the household survey results. The two estimates can then be compared.

To get an idea of the answer, estimates of the head-count index of poverty in each of 22 countries were first made from household survey data, using the WDR 1990 poverty line of $370 per capita, at purchasing power parity in 1985.[52] For the purpose of this study, available aggregate data were then used to predict the level of poverty in each country. This was done using a regression model for each country calibrated to the data for the other 21 countries.[53] Thus, 22 separate regressions were estimated i.e., the estimate for any one country does not assume that one knows the poverty measure for that country particular country, though one knows it for each of the other 21 countries. The predictor variables were private consumption per capita from the national accounts evaluated at purchasing power parity as well as official exchange rates, the level of urbanization, the infant mortality rate, life expectancy at birth, and the proportion of the labor force who are women. (A number of other social indicators were also tried, but did not help improve the measure of predictor precision.) It is unlikely that better predictions could be made from the available social and economic indicators. So the results presented here will probably err on the side of over-estimating the precision with which the prevalence of poverty can be estimated without a household survey.

---

52 These were the poverty estimates used in Chapter 2 of the World Development Report 1990; the methodology is documented in Ravallion, Datt and van de Walle (1991).

53 The modelling procedure adopted was the same as that described in Ravallion, Datt and van de Walle (1991) for extrapolating from countries where distributional data were available to countries where it was not.

Figure 6: Estimating the Head-Count Index Without a Household Survey

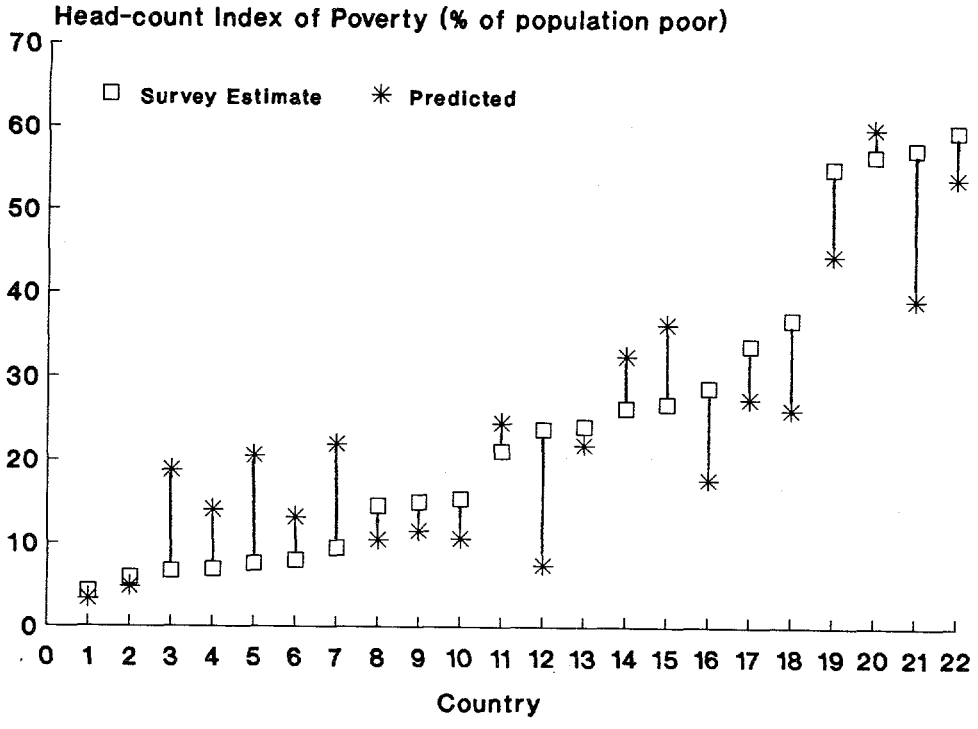

Note: Estimates of poverty in various countries using household survey data as compared to the results obtained by extrapolation using aggregate indicators and poverty measures for other countries.

Source: author's calculations.

Figure 6 summarizes the results. The countries have been ranked in ascending order of the estimates obtained from the household surveys. The predictions from the aggregate data are then plotted for each country.

Though the two estimates are positively correlated (the simple correlation coefficient is 0.87), it can be seen from the figure that the absolute errors in predicting the prevalence of poverty are often quite large. Indeed, the average absolute error as a percentage of the original survey estimate is 49%. The poorest few countries were correctly identified, though there are still a lot of changes in ranking amongst the countries. For example, the aggregate indicators suggest

that country 12 (the 11th poorest from the survey estimates) is almost the least poor amongst the 22 countries. The largest absolute error is country 21, for which the survey estimate of the head-count index is 57%, while the aggregate economic and social indicators predict a figure of 39%.

These discrepancies arise from two distinct factors: First, the extent of inequality varies from country to country, and this is hard to pick up well without distributional data from a household survey. Second, countries differ in the relationship between the social indicators and the consumption based poverty measures; some countries where consumption poverty is high have quite good social indicators, such as low infant mortality due to effective public health care, while others do not. These differences amongst countries can make it hard to assess the extent of poverty without a survey. The readily available economic and social aggregates can give, at best, a rough idea of the prevalence of poverty in a country.

3.2 How Well do Cross-sectional Indicators Identify the Long-term Poor?

The indicators of individual living standards discussed in section 2.2 and 2.3 which have been used in both poverty measurement and targeted poverty alleviation schemes typically measure the characteristics of households at a single date, or over a fairly short period. However, given that households face variability over time and can adapt to it, cross-sectional indicators may also reveal relevant information about circumstances at other dates, as discussed in section 2.2. This is important, as the alleviation of chronic poverty is clearly an important objective of redistributive policy.

Unfortunately, neither a priori reasoning nor the limited available evidence offer much guidance on the choice of a static indicator for identifying the chronically poor. For example, even when

households do successfully smooth their consumption over time, other factors such as measurement errors and the extent to which different static indicators move synchronously across households have bearing on the choice between current consumption and current income as indicators of chronic poverty.

Chaudhuri and Ravallion (1991) investigate what can be learnt from longitudinal data for guiding the choice of a static indicator in the more typical situations in which such data are unavailable. They propose a simple normative measure of the performance of static indicators in identifying the chronically poor. Performance is measured by the cost of the transfers needed to achieve any given impact on chronic poverty, when those transfers are made in a step-wise fashion starting from the person deemed to be poorest. They use this measure to investigate how well the most commonly used cross-sectional indicators do in identifying chronically poor households in three villages in dry-land areas of rural India.

The authors consider two measures of chronic poverty, one based on an eight year mean of income, the other based on a six year mean of consumption (the number of years used being determined by data availability). Six possible cross-sectional indicators of chronic poverty are investigated: current income, current consumption (net of durables and ceremonial expenses), current consumption (including these), current food consumption, the share of consumption going to food, and access to land.

They find that cross-sectional observations of consumption or income can correctly identify roughly three-quarters of the chronically poor in terms of either long-term income or consumption. However, this statistic can hide a great deal about the performance of the various indicators, mostly because it makes no distinction between identifying the poorest of the poor, and the not-so-poor.

Using their preferred measure of performance based on the Foster-Greer-Thorbecke P_2 measure, Chaudhuri and Ravallion find that current consumption is not a significantly better indicator of chronic poverty than current income. Indeed, current income is unambiguously the preferred indicator of chronic poverty based on mean income. The outcome does depend on which measure of chronic poverty is used; current consumption is better - though the gain is small - at revealing chronic poverty in terms of mean consumption at all except small budgets. Figure 7 shows how much impact on chronic poverty is obtained using various indicators for various budgets. The highest point reached on the vertical axis gives the initial poverty level, and each point on a given curve gives the final poverty level (vertical axis) which can be reached when a given budget (horizontal axis) is allocated according to the cross-sectional indicator corresponding to that curve. Thus, the lower the curve at any given budget the better the indicator is in reducing poverty with that budget. While income is generally the better indicator, there is not much to choose between it and consumption. And it should also be noted that the comparison here is solely in terms of the ability of these indicators to identify the chronically poor; it can be argued that current consumption is still the better indicator of poverty at a given date.

Two other indicators that have been used in both research and policy are found to perform rather badly, namely access to land and (especially) the share of food in consumption (Figure 7). Indeed, it would be better to simply give everyone (whether identified as poor or not) the same amount than base targeting on the food share. In some years, you could do better in identifying the chronically poor by making a purely random choice than by using food-share (or, better still, pick households with low food-shares!). This reflects the fact that the income elasticity of demand for food is close to unity in these

Figure 7: Effects on Chronic Poverty of Targeting Using Various Static Welfare Indicators

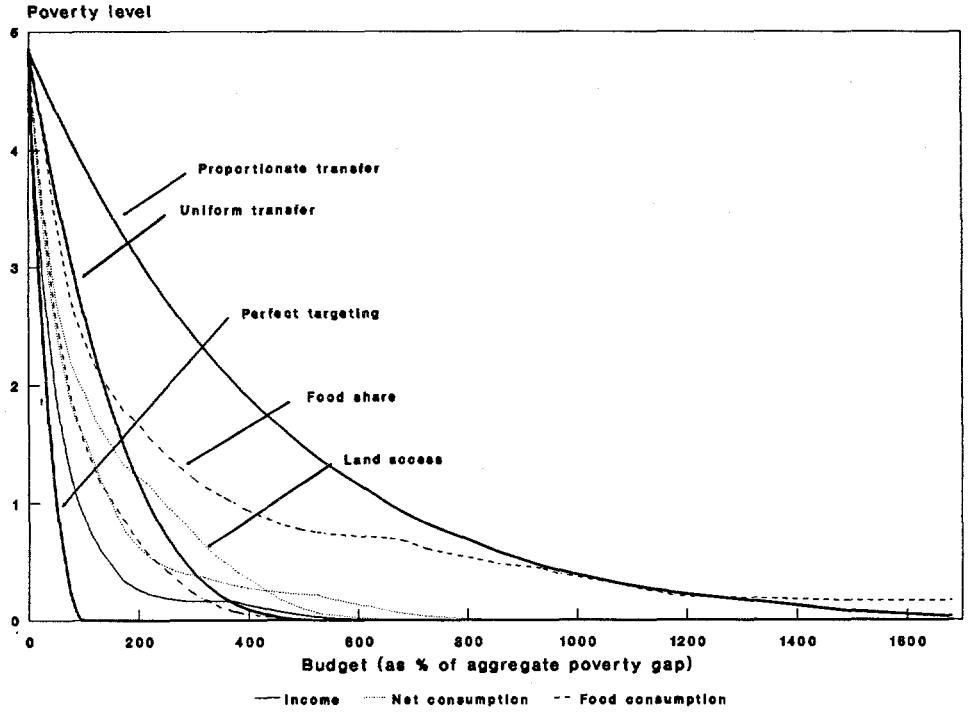

Note: The figure shows the rate of decline in chronic poverty (as indicated by eight year mean income) when step-wise targeting (from the bottom up) is based on various cross-sectional indicators for three villages in India.

Source: Chaudhuri and Ravallion 1991.

villages. The poor performance of access to land casts doubt on the efficacy of the various forms of land-contingent targeting that have been popular with policy makers in the subcontinent. Section 3.7 will return to this point.

The study's results also throw some light on the cost-effectiveness of targeting based on static indicators, and the potential gains from longitudinal observations. Using either current consumption or current income, the authors find that one can easily halve the transfer cost of a moderate impact on chronic poverty that would be

incurred without targeting. Nonetheless, the cost could still be as much as twice that under perfect targeting using longitudinal data. These estimates do not, however, incorporate any differences between indicators in the costs of data collection and transfer administration. Such costs would need to be considered before a policy conclusion can be drawn. The common use of poverty gap indices as measures of the cost of eliminating poverty through policy intervention is questionable in the usual situation in which one is constrained to using cross-sectional information. The current poverty gap may greatly underestimate that cost.

3.3 Which Sector or Region has More Poverty?

Obtaining a reliable poverty profile is an important step in applied poverty analysis, as it often directly informs policy discussion. If one is happy with the chosen welfare indicator, then it is largely a matter of selecting interesting partitions of the population into the sub-groups which constitute the poverty profile, then using a constant poverty line across all sub-groups. However, there are often important factors influencing welfare which are not adequately incorporated into the welfare indictor, such as spatial cost-of-living differences. One way to deal with these is to set a different poverty line across different sectors or regions. In this section I will first examine the problems encountered in one common method of setting poverty lines for different sectors. After that, examples of two detailed poverty profiles will be given, one regional, one sectoral. Both will also be used in later applications.

Urban-Rural Poverty Lines in Indonesia

The basic needs approach to constructing poverty lines (section 2.4) is the most common method used in developing countries. Section

2.4 mentioned the need for care when using this approach in constructing poverty profiles. This section will present an example.

Indonesia's Central Bureau of Statistics (Biro Pusat Statistik: BPS) uses a version of the "food energy method" (section 2.4) for constructing its poverty lines. It proceeds by first fixing a food energy intake cut-off in calories, and then finding the consumption expenditure at which a person typically attains that food energy intake. One then counts the number of people with consumption expenditure less than this amount. Thus one is estimating the number of people whose total consumption expenditures would be insufficient to attain the pre-determined food energy intake, given the prevailing relationship between food energy intake and total consumption across the population. The method is applied separately to each sector (urban/rural) and each date. The BPS method (or variations on it) has been used in poverty studies for other countries. The Indonesian practice is not un-usual.

Clearly one wants the poverty lines used to properly reflect differences in the cost-of-living across the sectors or dates being compared. However, as discussed in section 2.4, the food energy method is quite unlikely to generate poverty lines which are constant in terms of real consumption or income across the sectors/dates being compared. The reason is that the relationship between food energy intake and consumption or income is not going to be the same across sectors/dates, as noted in section 2.4. And there is nothing in the methodology to guarantee that these differences are ones which would be considered relevant to absolute poverty comparisons.

The specific example of the food energy method used by BPS generates differentials in the poverty line between urban and rural areas which appear to be far in excess of the cost-of-living differential. The differentials over time tend also to exceed the rate of inflation. As is typically the case in developing countries, the relationship between food energy consumption and total expenditures is

very different between urban and rural areas, with higher calorific intakes at any given consumption expenditure level in rural areas. This could well reflect the fact that agricultural work tends to be far more strenuous than most urban activities, and thus entails higher food-energy requirements to maintain body-weight.[54] Differences in relative prices and tastes may also be important. The relationship between calorie intake and income or consumption appears also to be shifting over time, with progressively lower food energy intakes at any given real income level.

The difference in the food-energy and income relationship between urban and rural areas is so large that, at any given food-energy requirement level, the urban poverty line exceeds the rural poverty line by a magnitude which is sufficient to cause a rank reversal in the estimated head-count index of poverty between the two sectors. This is illustrated in Figure 8, which gives the cumulative distribution of nominal consumption per person in each of urban and rural areas of Indonesia in 1987. At the BPS rural poverty line for 1987 of Rp. 10,294 per person per month, one finds that about 18% of the rural population is poor. At the BPS urban poverty line of Rp. 17,381 we find that 20% of the urban population is poor. But notice that, at any given poverty line (fixed across both sectors), the proportion of the rural population deemed poor is <u>higher</u> than that of the urban population. And this holds wherever one draws that poverty line.

However, the distributions in Figure 8 are not adjusted for cost-of-living differences. Table 5 gives my estimates of the critical poverty line differential for a reversal of the sector poverty ranking, using the 1984 and 1987 distributions of consumption per capita from the SUSENAS. Thus, for 1987, as long as the urban poverty line is no more than 60% higher than the rural poverty line, the head-count index will

---

54 See, for example, the estimates of caloric requirements for various activities given in WHO (1985).

Figure 8: Urban-Rural Poverty Comparison for Indonesia

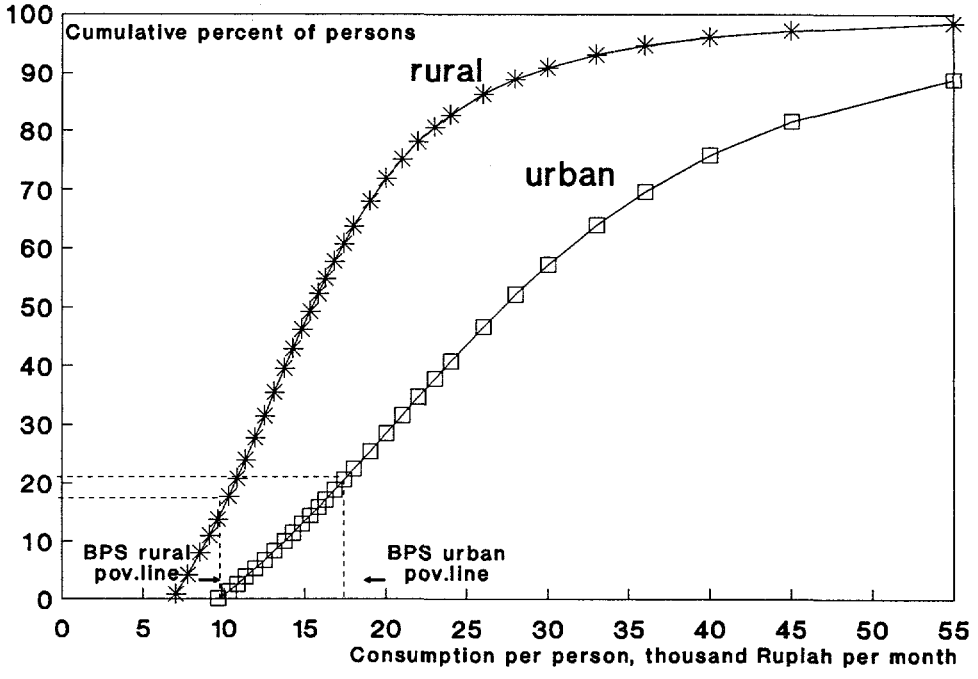

Note: Cumulative frequency distributions of nominal consumption per person in Indonesia 1987 for each of urban and rural areas.

Source: author's calculations from 1987 SUSENAS data tapes.

Table 5: Critical Cost-of-Living Differentials for a Reversal in Ranking of Poverty Measures Across Urban and Rural Sectors of Indonesia

	1984		1987	
Critical cost-of-living differential for rank reversal (urban/rural, %)*	71(U)	71(R)	62(U)	60(R)
BPS poverty line differential (urban/rural, %)	77		69	

Note: * (U) means that the BPS urban poverty line was used as the reference, while (R) used the BPS rural poverty line.

Source: Author's calculations from 1984 and 1987 SUSENAS data tapes.

be higher in rural areas. But with the BPS differential of 70% we get the opposite conclusion.

Unfortunately, there is no entirely satisfactory index for spatial cost-of-living comparisons in Indonesia. However, from what we know about price differentials between urban and rural areas, cost-of-living differences as high as 70% are quite implausible. Rice prices are generally about 10% higher in urban areas than rural areas, though in Jakarta they may be as much as 20% higher than in rural areas of Java. Housing costs are clearly much higher in urban areas, though this reflects in part quality differences. Ravallion and van de Walle (1991a) estimated a behavioral cost-of-living index for Java using a demand model estimated on 1981 data, which allowed for housing cost differences (after controlling for observable differences in housing quality), and rice price differences. For the poor, the estimated cost-of-living difference between urban and rural areas was about 10%, though it was slightly over 20% between Jakarta and rural areas. Furthermore, from casual observations one would suspect that there are some manufactured goods which are actually cheaper in urban areas. Though the evidence is still incomplete, the most plausible conclusion is that, at the likely cost-of-living difference between urban and rural areas, the incidence of poverty in Indonesia is higher in rural areas than urban areas. The food energy method of calculating the poverty lines suggests the opposite conclusion.

These problems can become quite serious when there is <u>mobility</u> across the groups being considered in the poverty profile, such as migration from rural to urban areas. For suppose someone who is just above the poverty line in the rural sector moves to the urban sector and obtains a job there generating a real income gain less than the difference in poverty lines across the two sectors. Though that person is better off in terms of the indicator of living standards being used, the aggregate measure of poverty across the sectors will show an

increase, as the migrant will now be deemed poor in the urban sector. Indeed, it is possible that a process of economic development through urban sector enlargement, in which none of the poor are any worse off, and at least some are better off, would result in a measured increase in poverty. With the likely importance of urban sector enlargement in the future evolution of poverty in Indonesia, this could become a quite misleading property of the BPS poverty measures.

The key analytical point here is that (in this context) one should be wary of making poverty comparisons in which the poverty lines used do not have a constant value in terms of the individual indicator being used in deciding who is poor. If one is measuring living standards using household consumption per person, adjusted for differences in the cost-of-living, then I would recommend that the poverty line should have constant value in terms of this indicator across the groups of households, or dates, being compared. Alternatively, if one is aiming to measure the more narrow concept of "food-energy poverty" or "undernutrition" then one must be careful that appropriate caloric requirements are used for each sector, reflecting the activity levels typical of those sectors.

Examples of More Detailed Sectoral and Regional Poverty Profiles

Table 6 gives an example of a poverty profile in which the sampled households in Indonesia's 1987 SUSENAS have been classified into 10 groups according to their principal income source. This is a summary of a more detailed sectoral poverty profile given in Huppi and Ravallion (1991). Results are given for the three poverty measures discussed in section 2.5. The following points should be noted:

i) An urban-rural cost-of-living differential of 10 percent has been assumed. From the only study that has been done this appears to be a reasonable assumption (Ravallion and van de Walle 1991), though it is

a good deal lower than has been assumed in other poverty profiles for Indonesia, as discussed above.

ii) 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.

iii) 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 un-wieldy poverty profile. An alternative is to calculate average incomes from various sources for various consumption groups ("ultra-poor", "poor", "near-poor", "others", say); see, for example, Huppi and Ravallion (1991).

iv) The three measures are in close agreement on the ranking of sectors in terms of poverty, with very few re-rankings. For example, the two farming sub-groups are the poorest for all three measures.

Table 7 gives an example of a regional poverty profile. The same three poverty measures are given for the main states of India based on the 1983 NSS. Points to note are that:

i) The methodology is similar to that used in constructing the Indonesian poverty profile discussed above. The distributions are of persons ranked by consumption per person of their household. Each state distribution was constructed from the underlying urban and rural distributions with an allowance for urban-rural and inter-state cost-of-living differences, as discussed in Datt and Ravallion (1991a).

Table 6: A Sectoral Poverty Profile for Indonesia in 1987

Principal Sector of Employment	Population share (1987)	Poverty Measure: Head-count index (H)	Poverty gap index (PG)	FGT P_2 measure
Farming				
(Self-emp.)	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).

ii) Results are given for two poverty lines commonly used in the literature on poverty in India (Datt and Ravallion 1991a).

iii) Again one finds strong agreement on the ranking of states according to the three poverty measures and both poverty lines. For example, the same four states - Bihar, Orissa, West Bengal and Tamil Nadu - are ranked in the first four for all three poverty measures and both poverty lines.

Table 7: A Regional Poverty Profile for India, 1983

State	Lower Poverty Line (Rs 77/ps/mn)			Upper Poverty Line (Rs 89/ps/mn)		
	Poverty Measure					
	H	PG	P_2	H	PG	P_2
Andhra Pradesh	20.50	4.34	1.43	30.44	7.27	2.56
Assam	25.82	4.23	1.07	41.58	8.32	2.41
Bihar	47.98	12.71	4.62	60.76	18.52	7.48
Gujarat	29.31	5.55	1.55	43.04	9.81	3.13
Haryana	12.89	2.26	0.65	21.69	4.33	1.32
Himachal Pradesh	17.36	2.81	0.72	28.85	5.62	1.62
Jammu and Kashmir	9.73	1.32	0.29	19.56	3.14	0.78
Karnataka	34.81	9.31	3.47	45.11	13.57	5.54
Kerala	27.87	6.18	2.01	38.82	9.96	3.58
Madhya pradesh	30.44	6.81	2.17	41.78	10.88	3.90
Maharashtra	35.69	9.18	3.28	46.65	13.62	5.38
Manipur	17.39	3.04	1.19	30.93	5.94	2.02
Meghalaya	29.59	8.30	3.19	38.44	11.87	4.98
Orissa	42.69	11.50	4.41	55.16	16.71	6.93
Punjab	11.62	2.09	0.60	19.35	3.94	1.21
Rajasthan	24.06	5.51	1.86	33.58	8.75	3.22
Tamil	40.74	11.59	4.63	51.63	16.41	7.06
Tripura	23.06	4.33	1.24	35.01	7.76	2.47
Uttar Pradesh	30.15	6.85	2.25	41.48	10.88	3.96
West Bengal	43.31	13.04	5.51	54.37	18.02	8.11
Total	32.65	8.09	2.90	43.90	12.29	4.79

Note: All poverty measures are expressed as percentages.

3.4 How Reliable are Assessments of Progress in Poverty Reduction?

Given the many uncertainties in measurement (due to errors in the survey data, and other sources, such as in setting the poverty line), it is important to know how sensitive conclusions about progress in poverty alleviation are to changes in the measures used. Here are brief descriptions of two case studies in testing the robustness of poverty comparisons over time. The methods used are easily applied to the sort of data typically available for country poverty assessments.

Bangladesh in the 1980s

The Bangladesh Bureau of Statistics (BBS) has recently published estimates of poverty for various years in the 1980s. Using the Household Expenditure Surveys (HES), BBS gives estimates of the absolute number, and proportion of the population, that were unable to attain a caloric intake of at least 2122 calories per person per day. These estimates are based on a comparison of actual household expenditure with an estimated expenditure needed to reach this caloric requirement, allowing for non-food consumption. Thus the BBS method of setting the poverty line finds the consumption expenditure level at which a household typically attains the stipulated food energy requirement. This is done separately for urban and rural areas. (This also raises the problem discussed in the previous section, but we leave that problem now.)

The BBS results suggest a substantial decrease in the numbers of persons deemed poor during the 1980s, with a decrease in the head-count index from 71% to 36% over the four years from 1981/82 to 1985/86, implying that aggregate numbers of poor declined from 67 million to 51 million. If one believes these figures then the performance is impressive indeed.

However, skeptics have pointed to a number of possible problems in comparing BBS's Household Expenditure Surveys over time. There are differences in sample size; while the 1981/82 HES surveyed about 9,500 households, this dropped to just 3,800 for the 1985/86 HES. However, a smaller sample need not be less reliable. More worrying is that there were also some changes in the questionnaire used, with implications for the measurement of household incomes and consumptions, such as in the methods used for imputing the value of food consumption from own production. Incentives to respondents (free blankets were apparently given) were also cut in the later survey, with likely implications for the response rate of the poor.

So there are good reasons to ask: How robust is this seemingly excellent recent record on poverty alleviation to possible inconsistencies over time in the underlying household expenditure surveys, on which the estimates of poverty have been based?

One possible clue can be obtained from an alternative source of data on aggregate incomes and consumption, namely Bangladesh's national accounts (NA). It cannot be assumed that the national accounts are more accurate, but dramatic discrepancies with that source would clearly be worrying. The national accounts may not provide a better estimate of average consumption than the HES, but (given the aforementioned changes over time in HES methodology), the NA series is probably a better basis for comparing consumption aggregates during the 1980s.

National accounts and household expenditure figures often give quite different estimates of average consumption. One is never sure which is more accurate; the national accounts typically estimate consumption as a residual, and other errors tend to end up in this category. On the other hand, household surveys plausibly under-estimate consumption, particularly of luxury goods by the rich (though this is not, of course, a problem for poverty measurement). However, for poverty comparisons it is more worrying if the two sources give very different estimates of the rate of growth in average consumption, since this is an important factor in assessing how absolute poverty has evolved over time.

The rate of growth in real consumption per capita implied by the HES is 10% per year over this four year period, while the national accounts suggest a much lower rate, about 0.5% per year. Such a large difference in growth rates is clearly worrying. It is consistent with the claims that the HES comparisons have over-estimated the rate of growth in consumption over this period.

How would this likely overestimation of the growth rate in real consumption affect the estimates of poverty in Bangladesh? To help

answer this question, I have calculated measures of poverty by two methods, one relying exclusively on the HES (following conventional practice), the other using information on mean consumption from the NA to supplement the data on relative inequalities from HES (Ravallion 1990a). The latter method also illustrates a potentially useful methodology for other purposes, so some elaboration is called for. The idea is that one derives formulae for the poverty measures of interest as functions of the mean of the distribution, and a set of other parameters describing the Lorenz curve. The latter parameters are estimated econometrically. From these formulae one can then estimate the poverty measure that would be obtained if the mean changed, holding the Lorenz curve constant; thus one can estimate the poverty levels that would hold if the mean was that obtained from NA, rather than HES, holding the Lorenz curve constant.

We should be clear about the purpose of these calculations. Their aim is not to come up with the "best" measure of the magnitude of poverty in Bangladesh. Rather, the objective is to test the robustness of the BBS estimates of how poverty has changed in Bangladesh during the 1980s.

The results are summarized in Table 8, which gives both the head-count index, and the two other poverty measures in the FGT class, PG and P_2. Estimates implied by the HES and BBS methodology are compared with those implied by the HES Lorenz curve and the growth rate from the national accounts when applied to the HES mean for 1981/82.

A number of points emerge. In marked contrast to the HES/BBS estimates, the assessments of growth consistent with national accounts data (only using the HES to measure <u>relative</u> inequalities) suggest that the proportion of the population deemed to be poor has remained fairly stable over the period. Absolute numbers of poor increased. It is clear that the conclusion that poverty decreased in Bangladesh during this period must be viewed with some skepticism.

Table 8: Alternative Measures of Poverty in Bangladesh

Poverty measure		1981/82	1983/84	1985/86
Head-count index (H)	HES	71.2	43.7	35.9
	NA	71.2	67.9	70.5
Poverty gap index (PG)	HES	24.2	11.5	7.8
	NA	24.2	22.0	22.9
Foster-Greer-Thorbecke (P_2)	HES	10.6	4.4	2.5
	NA	10.6	9.5	9.5

Note: All poverty measures expressed as percentages.
HES: Sample mean from Household Expenditure Survey
NA: Mean consistent with growth rates from national accounts, applied to 1981/82 mean from HES.

Source: Ravallion (1990a).

Indonesia in the 1980s

Ravallion and Huppi (1991) look at the evolution of aggregate poverty in Indonesia between 1984 and 1987, a period embracing external shocks and macroeconomic adjustments. They examine a number of welfare indicators: real consumption per person, real income per person, food energy intake per person, and the budget share devoted to food. These were calculated from the data tapes of the 1984 and 1987 SUSENAS, with 50,000 households randomly sampled at each date. (The samples were stratified, and the distributions were corrected for this. The Consumer Price Index was also re-weighted to reflect the consumption pattern of the poor.) For all indicators, Ravallion and Huppi found that first-order dominance holds over the period. Figure 9 gives the first-order dominance test: the 1984 distribution of consumption per capita lies entirely <u>above</u> 1987. Any poverty line (even up to the highest consumption level!), and any reasonable poverty measure will give the same conclusion: poverty fell. The other welfare indicators show a similar result.

Figure 9: Poverty Incidence Curves

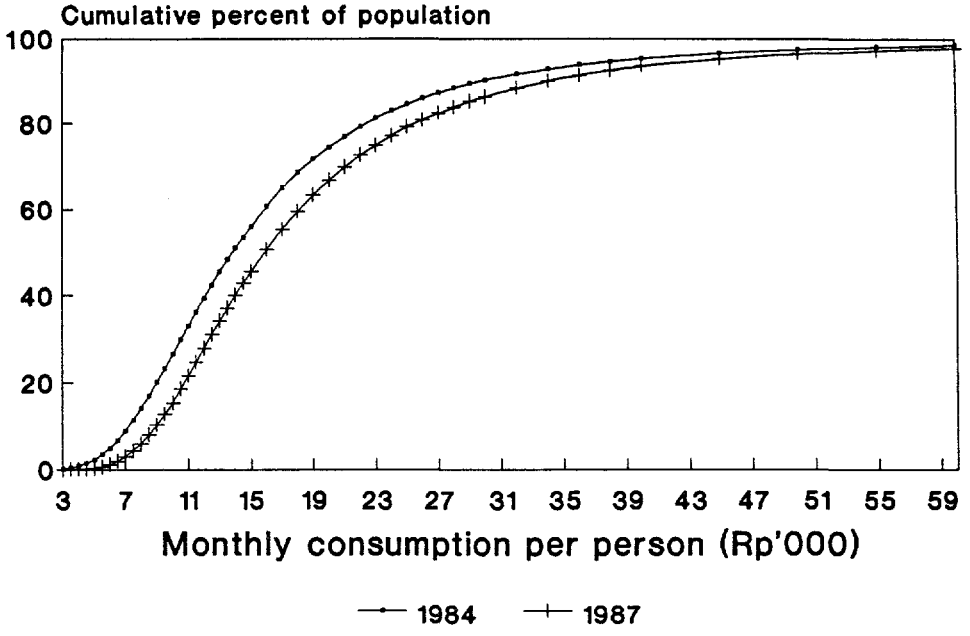

Note: First-order dominance test based on cumulative frequency distributions of consumption per capita in Indonesia, 1984 and 1987.

Source: Ravallion and Huppi 1991.

Furthermore, unlike the results discussed above for Bangladesh, the conclusion that poverty fell appears to be quite robust to likely measurement error in the underlying household survey, and the price index. For example, Ravallion and Huppi provide similar calculation to those reported in Table 8, comparing the rates of growth implied by the national accounts with those from the household surveys. They also test for sensitivity of their first-order dominance conclusion to possible under-estimation of the rate of inflation over the study period; they find that the actual rate of inflation would have to have been 50% higher than that recorded to alter the dominance result - this is an unlikely magnitude of error. Despite various measurement problems, the authors conclude that the qualitative result for Indonesia as a whole is robust: poverty fell.

Given the regional diversity of Indonesia it is of interest to also ask how the regional poverty profile changed. Again the problems in adjusting for spatial cost-of-living differences loom large in attempting to construct regional poverty profiles using real consumption as the welfare indicator. However, comparisons of the more narrow concept of undernutrition, based on the food energy values of consumption, do not require price data. Ravallion (1992a) has constructed distributions of food energy intakes for each of urban and rural areas of each of Indonesia's provinces for 1984 and 1987 giving 52 regions. The theory of stochastic dominance can be used to test for a decrease in undernutrition even when the distribution of food energy requirements is unknown (section 2.7).

The first-order dominance test indicates an unambiguous improvement over the period for 29 of the 52 regions, while an unambiguous worsening is indicated for only three (Ravallion 1992a). However, the first-order test was ambiguous for 20 regions. For each of these, the 1987 distribution crossed the 1984 just once from below i.e., there was an improvement at the lower end of the distribution. The average cross-over point was at about 1950 calories per person per day. The second-order test resolved this ambiguity for 12 regions, all but one of which showed an unambiguous improvement.

Clearly then there was unevenness across regions in progress in reducing undernutrition, though first-order or second-order dominance tests do indicate a reduction in undernutrition for about 80 percent of the regions. The quantitative extent of this improvement also varies considerably across regions; see Ravallion (1992a) for an analysis of the factors explaining these differences.

3.5 What is the Relative Importance of Growth Versus Redistribution?

Datt and Ravallion (1991b) illustrate the decomposition discussed in section 2.6 with a comparative analysis of the recent evolution of poverty measures for Brazil and India.

India's performance in poverty alleviation was better than Brazil's over the difficult 1980s. For example, toward the end of the 1980s, the two countries had an almost identical poverty gap index PG (though Brazil's being for a local poverty line with higher purchasing power), while India's index of 10 years earlier had been about 50% higher than Brazil's had been at the beginning of the 1980s. India's progress over the 1980s has been uneven across sectors, with the urban sector contributing a rising share of aggregate poverty.

In comparing these countries, their results indicate quite different impacts on the poor of distributional changes over the 1980s. Over the longest periods considered by Datt and Ravallion, distributional shifts have aided poverty alleviation in India at a given mean consumption, while they have hindered it in Brazil. Without any change in the mean, India's PG index would still have fallen quite noticeably (from 16% of the poverty line to 11% in rural areas) while Brazil's would have increased equally sharply (from 10% to 13%). With Brazil's worsening distribution (from the point of view of the poor), far higher growth rates than those of the 1980s would have been needed to achieve the same impact on poverty as India attained.

Growth and distributional effects on poverty were quite uneven over time in both countries, and the effects of instances of negative growth were notably different between the two. Contraction in the mean due to the poor agricultural years of 1986 and 1987 was associated with a (modest) improvement in distribution in India, such that poverty continued to fall (at least by the distribution sensitive P_2 measure). Contraction in Brazil due to the macroeconomic shocks of the 1980s was

associated with a marked worsening in distribution, exacerbating the adverse effect on poverty.

3.6 How Important are Different Sectors to Changes in Poverty?

Table 9 provides information on the relative contribution of various sectors to aggregate poverty alleviation in Indonesia over the period 1984 to 1987. These are the "intra-sectoral effects" in equation (13), expressed as a percentage of the reduction in aggregate poverty for each poverty measure. 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. The second most important contribution came from gains to farm workers, whose reduction in poverty as measured by the head-count index contributed 11% to the reduction in the aggregate index, while the decline in this sector's P_2 measure contributed almost 17% of the aggregate decline. These two groups jointly accounted for 61% of the reduction of the aggregate head-count index and 74% of the reduction of the aggregate value of 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.

Table 9: Sectoral Decomposition of the Change in Poverty in Indonesia, 1984-1987

Principle Sector of Employment	Population share (1984)	Contribution of sector to change in: Head-count index (H)	Poverty gap index (PG)	FGT P_2 measure
Farming				
(Self-emp.)	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 sectoral 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

Note: minor sectors omitted.

Source: Huppi and Ravallion (1991).

Other sectors contributed less to poverty reduction during the period, both in absolute terms and (generally) relative to their population shares. Gains within the rural non-farm sectors (industry, trade and services) were larger than the gains attributable to the corresponding urban sectors.

Also noteworthy is the relatively important part of aggregate poverty reduction due to population shifts.[55] Over 13% of the decline in the national head-count index was due to population shifts between various sectors of employment, and over 9% of the decline in the P_2 measure can be traced back to these shifts. The sectors which gained in population share were almost all urban (Huppi and Ravallion 1991), and had initially lower poverty measures. This is the main factor underlying the contribution of population shifts to poverty alleviation. The fact that population was moving out of the rural sector, where poverty was falling faster, accounts for the negative interaction effects in Table 9.

3.7 How do Price Changes Affect the Poor?

One of the main routes through which macroeconomic policy changes affect the poor is through shifts in relative prices. However the poor tend to be heterogeneous in terms of their net trading position in many relevant markets, such as for food staples; while some are net producers others are net consumers. This section will use some case studies to show how the methods of both welfare measurement and poverty measurement discussed in section 2 can help assess impacts on the poor of price changes.

Rice Prices and Poverty in Indonesia

How does an increase in the relative price of a food-staple affect poverty? To answer this question convincingly we need a welfare measure which reflects the effects of price changes in a way which is consistent with demand and supply behavior at the household level. The <u>equivalent income</u> concept can be used here, a generalization of the idea of real income (money income per equivalent adult deflated by a price index).

---

55 Though urban-rural migration is not the only cause; there was a certain amount of re-classification of rural areas as urban areas.

To outline the idea of equivalent income properly we need to use modern consumer theory. The following exposition will give an intuitive idea of what is involved.

Let $v(p, y, x)$ denote the level of welfare of a household facing prices p, having income y (which may also be a function of p), and having other relevant characteristics x. Consider some <u>reference values</u> for p and x which are fixed across all households. Then the <u>equivalent income</u> (y^e) can be defined implicitly by equating:

$$v(p^r, y^e, x^r) = v(p, y, x) \tag{14}$$

When so defined, the equivalent income is an exact monetary measure of the household's welfare; the equivalent income is the income (or consumption expenditure) level which a given household would need to achieve their actual level of welfare in the hypothetical situation in which that household faced the same prices as all other households, and had the same size and demographic composition. The concept thus allows the analyst to retain the appeal of using a monetary measure of well-being, while adjusting that measure for differences in prices and household characteristics.

By first estimating a suitable (including utility consistent) demand model one can estimate the equivalent income at various values of p, y, and x, for the fixed reference variables. A policy reform can be interpreted as a change in (p, y), and one can then trace out the effects on welfare. Following the discussion in sections 2.2 and 2.3, the analyst should be aware of the limitations of this exercise. In practice, there will always be aspects of individual well-being which cannot be identified by looking at demand behavior, and so other information on the impacts of the reform should not be ignored such as, for example, adverse effects on the nutritional status of children (section 2.3).

Figure 10: First-Order Dominance Test for a Change in Food-Staple Price

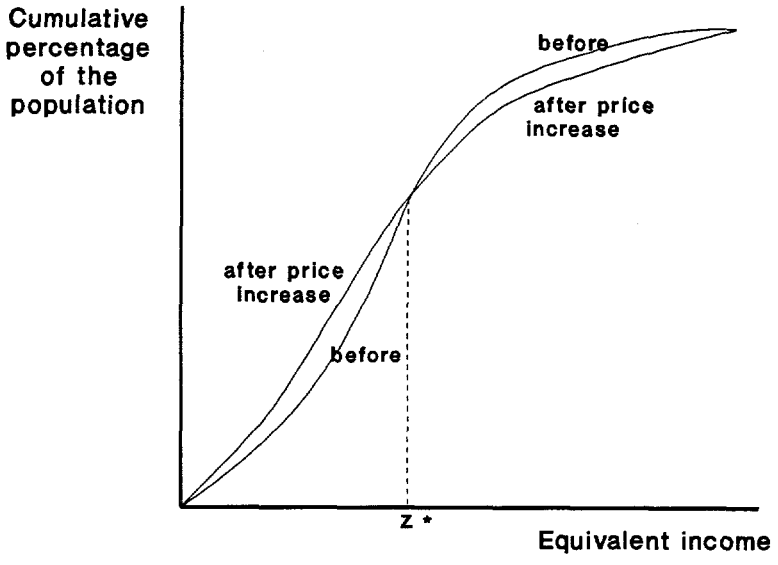

Note: The cumulative distributions of equivalent income before and after an increase in the price of rice.

Source: stylized, but consistent with results for Indonesia in Ravallion and van de Walle 1991.

Ravallion and van de Walle (1991a) used this approach in studying the effects of rice price changes on poverty in Indonesia. They begin by estimating the equivalent income function using a demand model, and then simulate the new values of each household's equivalent income after the stipulated rice price changes.[56]

They found that the choice of poverty line was crucial to assessments of the effect on aggregate poverty of an increase in the price of rice in the early 1980s, as implied by trade liberalization. Figure 10 shows how the cumulative distribution of equivalent income was found to shift after a rice price increase. The reason that first-order dominance fails in this case is that the poorest tended to be losers

56 Provided the price changes involved are not too large, this method can be simplified greatly by calculating a first-order welfare effect of the price change. An example will be given in the next section.

from a price increase (mainly landless rural poor), while many of the better off amongst the poor gained (rice surplus farmers). If the maximum possible poverty line is no higher than z* then the conclusion is clear: poverty will increase with the increase in the price of rice. The conclusion is ambiguous for higher values of z^{max}, and such values are not inconsistent with past poverty lines used for Indonesia.

However, the same study found that all measures of the depth of poverty show an unambiguous increase when rice prices increase, and this holds for all possible poverty lines. This is indicated by the fact that the area under the "before" curve in Figure 10 is everywhere below the area under the "after" curve. Second-order dominance permits a clear ranking. All poverty measures which increase (decrease) with a decrease (increase) in incomes of any of the poor will also increase (decrease) when the price of rice increases (decreases), and this holds for a wide range of poverty line.

Labor Market Responses to a Rice Price Change in Rural Bangladesh

The poor in rural economies of much of Asia, and increasingly in Sub-Saharan Africa, derive a large share of their income from wage employment in food-grain production. Higher rice prices will probably lead to higher agricultural wage rates. This will mitigate adverse effects on the poor. But how long will it take for labor markets to adjust? I have attempted to answer this question for Bangladesh (Ravallion, 1990b). The approach includes a dynamic econometric model of agricultural wage determination (Boyce and Ravallion, 1991). The wage rate at date t (w_t) was modelled as a function of lagged wages, current and past rice prices, and various other variables; in stylized form, the model can be written:

$$(15) \quad w_t = f(w_{t-1}, p_t, p_{t-1}, x_t)$$

This was then embedded in a farm-household model and used to estimate welfare effects on the poor of rice price changes.

The following formula can then be used to derive first-order estimates of the monetary value of the welfare effects of the price change at each income level:

$$(16) \quad \begin{array}{l}\text{monetary value to}\\ \text{farm household of}\\ \text{change in rice price}\end{array} = [\underbrace{p(Q^s-Q^d)}_{\substack{\text{value of}\\ \text{net supply}\\ \text{of rice}}} + \underbrace{w(L^s-L^d)E}_{\substack{\text{value of}\\ \text{net supply}\\ \text{of labor}}}] \times \begin{array}{l}\text{proportionate}\\ \text{change in rice}\\ \text{price}\end{array}$$

where E denotes the elasticity of agricultural wage rate with respect to the price of rice. In the special case in which E=0, this formulae collapses to the first-order analog of the method outlined in the previous section based on an estimated equivalent income function; the monetary value of the welfare change due to an increase in the price of some good, holding other prices constant, is approximated by the product of the change in price and the initial net supply (production - consumption) of that good. The approximation is better the smaller the price change and the less important are the general equilibrium effects on the prices of other goods. The results for Bangladesh yield E=.2 (short-run), E=.5 (long-run). These general equilibrium effects operating through the labor market turn out to have bearing on the welfare distributional effects of the rice price change, as discussed further in Ravallion (1990b).

The key implication from the study is that the rural non-poor tend to gain from higher rice prices in the short-run, while the rural and urban poor tend to lose. Welfare effects on the rural poor tend to be neutral or positive in the long-run. It would typically take 3-4 years before the rural poor will return to their previous welfare level.

3.8 What are the Gains to the Poor from Targeted Poverty Alleviation Schemes?

Directly targeted poverty alleviation schemes have been widely used in developing countries. Assessing impacts on the poor is an important but difficult analytical problem, recognizing that administrative capabilities typically fall well short of what would be needed for perfect targeting, and so some often subtle but real costs are incurred by both poor and non-poor in participating. This section will discuss some case studies on the assessment of such schemes.

Quantifying the Poverty Impacts of Regional Targeting

Administrative costs and related constraints on instruments for the direct alleviation of poverty are becoming more widely appreciated in analytical discussions of targeted policies. These constraints are particularly relevant in underdeveloped rural economies. We are nowhere near a point where one can meaningfully talk about negative income taxes in most settings; this is not a feasible policy option. The problems of observing incomes have led to a variety of schemes for "indicator targeting" whereby transfers are made contingent on correlates of poverty, such as landholding, caste, or place of residence.

Regional targeting of transfers has a number of obvious attractions when trying to reach the poor. Substantial regional disparities in living standards are common in developing countries, and backward areas can often be readily identified. Place of residence may thus be a useful indicator of poverty. Also, the existence of local governments suggests that an administrative apparatus is generally available, and has already been exploited by many poverty alleviation schemes in the developing world. For example, the allocation of central government disbursements across states in India has also been determined, in part, by inter-state disparities in the incidence of poverty.

Datt and Ravallion (1991a) have looked at this issue for India. They consider the effects on aggregate poverty of both additively and multiplicatively absorbed transfers across states of India, and between their urban and rural areas. Effects on pre-transfer incomes or relative prices are ignored. Depending on the assumptions made about how poverty is measured, they find that 75% and 90% of transfers in which the "donor" state or sector has a higher mean consumption level than the "recipient" would have reduced aggregate poverty in 1983. Thus, the qualitative effect of reducing regional/sectoral disparities in average living standards generally favors the poor.

However, the quantitative gains to the poor from even the complete elimination of regional disparities in average living standards in India are unlikely to be large. For example, Datt and Ravallion simulate the effects on various poverty measures of the complete elimination of regional disparities across 20 states of India, with each state divided into urban and rural areas, while holding intra-regional relative inequalities constant. This is done by first estimating each regional Lorenz curve, as discussed in section 2.5. The poverty measures are then derived as functions of these parameters and the mean of each local distribution. The effects of reducing regional disparities while holding intra-regional inequalities constant (ie., holding the Lorenz curve constant) can then be simulated numerically.

Table 10: Simulated Effects of Regional Equalization on Poverty in India

Poverty measure	Actual (1983 NSS, poverty line=Rs89/ pers/mnth)	Simulated with mean consumption equalized across all states and between urban and rural areas
Head-count index (H)	43.9	42.3
Poverty gap index (PG)	12.3	11.3
Foster-Greer-Thorbecke (P_2)	4.8	4.7

Note: Simulations assume that national mean is unaffected by regional re-distribution, and that Lorenz curves within regions are also unaffected.

Source: Datt and Ravallion (1991a).

The study found that regional equalization of average consumptions would only yield a small reduction in the proportion of persons below the poverty line, from an initial value of 44% to 42% after the elimination of regional disparities. The results are summarized in Table 10. Datt and Ravallion also examine how much extra impact on aggregate poverty is possible if information on the distribution within regions (from household survey data) is used optimally in regional targeting. Again they find that the potential impact is small.

The main reason for these results is that, despite the marked regional disparities in poverty measures revealed by Table 7, region of residence as defined by administrative boundaries in India turns out to be a fairly weak indicator of poverty; there are many non-poor within "poor" regions, and many poor within "rich" regions. The latter group bears part of the cost of regional targeting, while there is a leakage of benefits to the former group. Administratively feasible forms of regional/sectoral targeting may thus be quite blunt instruments for aggregate poverty alleviation in India. Growth costs of shifting

resources out of more profitable locations are likely to further reduce the gains to the poor.

Similar experiments for Indonesia, where the regional disparities in the incidence and severity of poverty across islands are larger, have suggested greater gains from this type of targeting - and also that those gains are far from being realized by existing inter-regional transfer policies (Ravallion, 1992b). But even then, regional targeting is no panacea; the impact on national poverty of unrestricted income redistributions across Indonesia's provinces would be equivalent to a three-four percent increase in national mean consumption. We should look for other indicators to enable finer targeting within regions or sectors.

Land-Contingent Targeting in Bangladesh

In rural areas of South Asia, and (increasingly) in sub-Saharan Africa, the most promising single additional indicator is probably land-holding class, given the strong negative correlation observed between land-holding and poverty in rural areas of much of South Asia. My simulations of the effects on poverty confirm the need for targeting poverty alleviation schemes toward the landless or near landless in rural areas of Bangladesh (Ravallion 1989). Figure 11 shows how the P_1 poverty measure varies with land-holding in rural Bangladesh, and also gives the transfers by land-holding class that would minimize the aggregate P_2 poverty measure for a given budget (see Ravallion 1989 for details). The transfer falls with land-holding, though post-transfer income still increases with land-holding even when poverty is minimized.

However, the results also highlight the limitations of land-contingent targeting. Though land-holding class is a strong correlate of poverty, it is not a perfect correlate; there are rural poor with ample land, and landless non-poor. Differences in land quality also add noise to the relationship. One clearly cannot expect to eliminate

Figure 11: Simulated Transfer Schemes

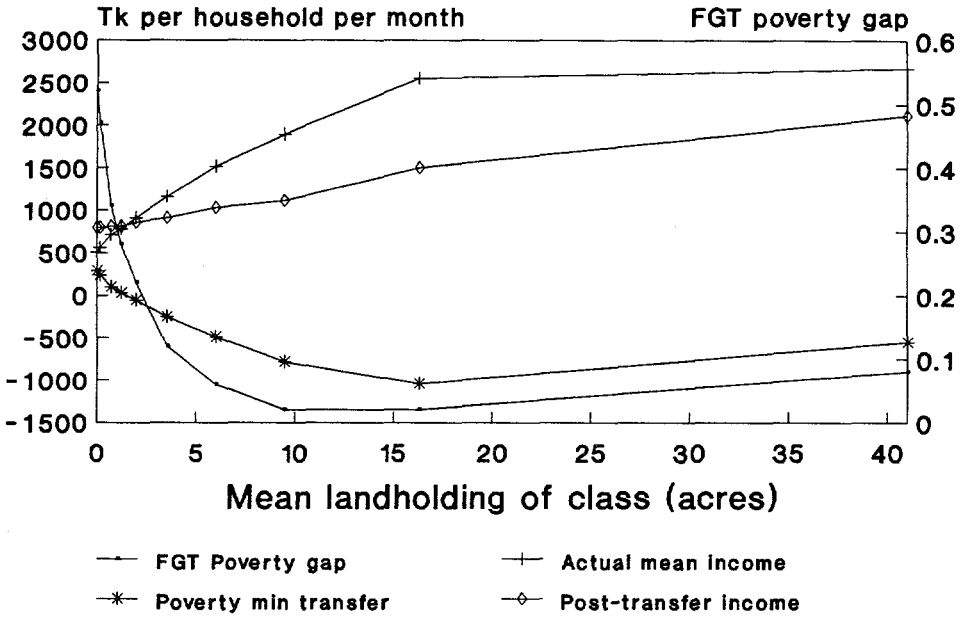

Note: Poverty minimizing allocation of land-contingent transfers in rural Bangladesh.

Source: adapted from Ravallion 1989.

poverty in this way, even with the most generous assumptions about what sort of income redistributions across land-holding classes are politically feasible. Indeed, even with complete control over the distribution of income across 10 landholding classes in Bangladesh, the maximum reduction in the aggregate severity of poverty which is attainable this way is no more than one could obtain by an un-targeted lump-sum gain to all households equivalent to about one tenth of GDP (Ravallion 1989). Various factors may enhance the poverty alleviation impact, such as any effects of the income or wealth gains on future productivity of the poor. Other factors will detract from their poverty alleviation impact. Plausible restrictions on the government's redistributive powers would further diminish the gains to the poor from such policies.

Public Employment Schemes

Disappointment with the prospects for direct poverty alleviation through other means, such as regional targeting, has re-kindled interest in one of the oldest known poverty alleviation schemes: relief work, such as the rural public works schemes common in South Asia, and becoming more common in other regions. The argument here is not that such work can alleviate poverty by creating income generating assets (though all the better if it does). Rather, it is that work requirements can provide seemingly excellent incentives for self-targeting in that the non-poor rarely want to participate and a great many of the poor do.[57]

How should these schemes be designed so as to be most cost-effective in alleviating poverty? The joint issue of the wage rate and coverage of a public employment scheme arises in policy design. For a given long-term budget, the choice is between a scheme which aims for wide coverage at potentially low wage rates, and one which rations participation at a wage rate sufficient to allow more participants to escape poverty.

I have examined this issue, and derived conditions for ranking stylized policy alternatives in terms of a broad class of poverty measures (Ravallion, 1991b). Empirical simulations of the policy alternatives for Bangladesh generally reinforce the case for wide coverage at low wage rates when the policy maker is concerned about the severity of aggregate poverty, such as measured by the Foster-Greer-Thorbecke measure P_2. This can hold even when the non-wage cost per worker is quite high (though if that cost is very high, limitations on coverage will become desirable). Wide coverage tends to reach the poorest first, and so it will have more impact on P_2 than on PG or H. A second-order dominance test generally favors wide coverage, so the conclusion is robust to the choice of poverty line and poverty measure,

57 For a survey of the theory and evidence see Ravallion (1991a).

provided the latter values reductions in the depth of poverty. Arguments for rationed coverage at higher wage rates are more convincing when one is only concerned about the incidence of poverty, as measured by H, since then one will find more people crossing the poverty line. The impact on poverty of a low wage public employment scheme is illustrated well by some results obtained by Datt and Ravallion (1991c). They studied the effects of public employment opportunities (mainly on Maharashtra's "Employment Guarantee Scheme") on incomes in a poor Maharashtran village using a longitudinal household level data set spanning six years. The data allowed estimation of the forgone incomes of the participants in the scheme, so that calculations could be made of the incomes that one would have expected if the scheme had not existed. Figure 12 gives the cumulative distribution of mean income per person, where the mean was formed over the six years of data, and can thus be considered a good indication of normal, "long-term", standard of living. Two distributions are given, with and without the public employment.

First order dominance must hold in this case; no incomes are any lower as a result of access to public employment.[58] However, it can also be seen that the scheme is targeted, in that the two distribution curves are not parallel. The largest vertical distance occurs just below Rs 700 per person per year, and is 13 percentage points; in other words, at a poverty line around this figure the head-count index of poverty would fall by this amount, from 33% to 20% as a result of the scheme.

However, typical poverty lines for rural India are well above this point. The most widely used poverty line would imply a figure around Rs 900 per person per year. At this point, the impact on the head-count index is much smaller; it drops from 42% to 40% as a result of the access to public employment. But, of course, this comparison hides the

---

58 The scheme is primarily financed by taxes on the urban sector of the state of Maharashtra.

Figure 12: Poverty Incidence Curves With and Without Earnings from Public Works in an Indian Village

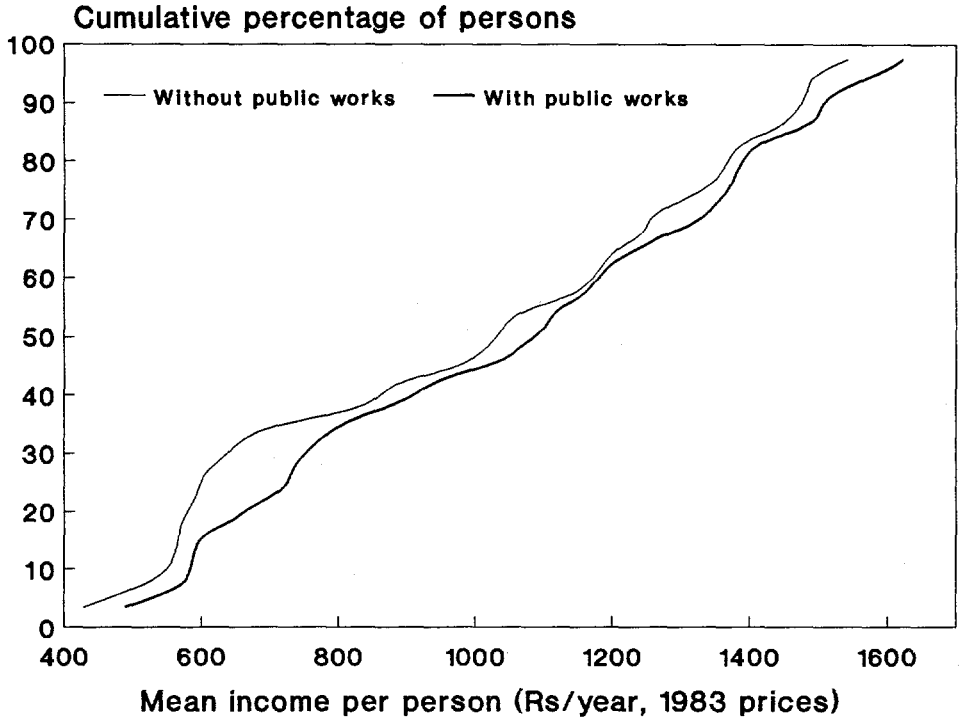

Note: Distribution of the six-year mean of household income per person in an Indian village, and the simulated distribution without access to public works employment, allowing for incomes forgone in alternative activities.

Source: Datt and Ravallion 1991c.

real impact on those well below the poverty line. A better indication can be obtained using a measure of the severity of poverty such as P_2; at the same poverty line, this measure shows a substantial decline as a result of the scheme, from 5.0% to 3.2%.

Public employment schemes thus illustrate how a policy assessment may rest heavily on the value judgments made in defining and measuring poverty. My own judgement - and one that I suspect many readers will share - is that policy assessments should not only be more sensitive to impacts on the "poor" than the "non-poor", but should also put highest weight on impacts on the poorest.

4. Recommendations

Assessment of development progress, and of the merits of specific policies, have often been informed by quantitative estimates of the magnitude, depth, and severity of absolute poverty. As we have seen a number of times in this monograph, such assessments can be quite sensitive to the measurement assumptions being made; for example, it is common for developing country poverty lines to be at or near the mode of the consumption distribution, at which point the head-count index of poverty will be very sensitive to the precise location of that line, since this is the point where the cumulative frequency distribution is steepest.

Fortunately, in many policy-oriented applications we are not so concerned about the precise estimate of a poverty measure, but rather we want to know which of two situations, such as before and after some policy change, has more poverty. I have tried to show how such poverty comparisons can be made, and what pitfalls are encountered in doing so. This section will distill those observations into some recommendations for best practice when making poverty comparisons.

In measuring living standards, the best indicator which is typically available is a suitably normalized and comprehensive measure of household consumption expenditure based on interviewee recall in a single visit household survey covering a representative sample. From the point of view of assessing living standards, one can do better than this, such as by using multiple interviews over time at the individual level. But such surveys are more costly for a given sample size, and so one faces a trade-off. The choice will depend on the priorities attached to the multiple purposes of survey data and the resources available. In years to come, I suspect that national level surveys will remain largely cross-sectional in developing countries, with longitudinal surveys typically being of far narrower coverage over the population, and being used for well focused questions, such as

investigating the dynamics of poverty at the individual or household level and in evaluating the impacts of specific policy interventions, possibly on an experimental basis.

Supplementary measures to those obtainable from household surveys are often useful, given the problems in finding the ideal monetary measure of well-being. Examples include the various "social indicators" (life expectancy, infant mortality, literacy), and indicators of access to public services (schools, health clinics, drinking water). These measures are valuable to the extent that they capture aspects of household well-being that are not adequately reflected in consumption or income based measures.

However, while these measures can provide useful supplementary data in making poverty comparisons, one should be equally wary of relying exclusively on such indicators. They need not be very informative about the living standards of the poor. In making comparisons over time, they may also be overly sensitive to one or two narrow dimensions of well-being. For example, average consumption (as estimated from the national accounts) may increase while the poor gained nothing. Some social indicators can also be overly sensitive to even a limited degree of progress in primary health care, particularly when starting from a high mortality rate. While preventing the common infectious and parasitic diseases such as malaria is undeniable progress, there is a lot more to the elimination of poverty. Furthermore, though some of these aggregate indicators are probably influenced by the distribution of welfare, their usefulness as indicators of distribution is quite unclear. For example, they may not be as responsive to the well-being of the ultra-poor as we would like. And the low participation of this group in many policies ostensibly designed to help them has been of serious concern.

Unfortunately barely a third of the developing countries have a recent (say for the 1980s) nationally representative household income or

consumption survey. And there is a marked regional variation in the availability of such data. Consider, for example, the poverty estimates for the developing world made for the 1990 World Development Report (World Bank 1990b). These relied on household surveys for 22 countries (though covering three-quarters of the population of the developing countries) and extrapolations using readily available aggregate economic and social indicators for the remainder (Ravallion, Datt and van de Walle 1991). Those extrapolations were, of course, imprecise. One can calculate just how much imprecision in the point estimate of poverty for any region is attributable to the extrapolation errors, reflecting the lack of survey data. Figure 13 gives the point estimate for each region of the developing world, and the 95 percent confidence interval around that estimate attributable solely to the imprecision arising from the need to make such extrapolations. It can be seen that the interval varies widely, being the most narrow for East Asia, where virtually the entire population is covered by a recent sample survey, and the most wide for Sub-Saharan Africa, where about 90 percent of the population is not covered by a recent survey.[59] Indeed though the point estimate of the head-count index for Sub-Saharan Africa is almost as high as that for South Asia (47 percent and 51 percent respectively), the confidence interval for Sub-Saharan Africa is 19 percent to 76 percent, while that for South Asia is 50 to 53 percent. Furthermore, it is plausible that, while the incidence of absolute poverty is falling in South Asia, it may well be increasing or at best stagnant in Sub-Saharan Africa (World Bank 1990b). It is likely that, within the next 10 years, we will find that we know least about poverty in the part of the world where its incidence and depth is highest.

---

59 This will start to improve as results of the various surveys being done for the Social Dimensions of Adjustment in Sub-Saharan Africa project, based at the World Bank, become available. However, even then, this is not expected to include most of the larger countries, notably Nigeria, Sudan, and Ethiopia.

Figure 13: Regional Differences in the Precision with which the Head-Count Index can be Estimated

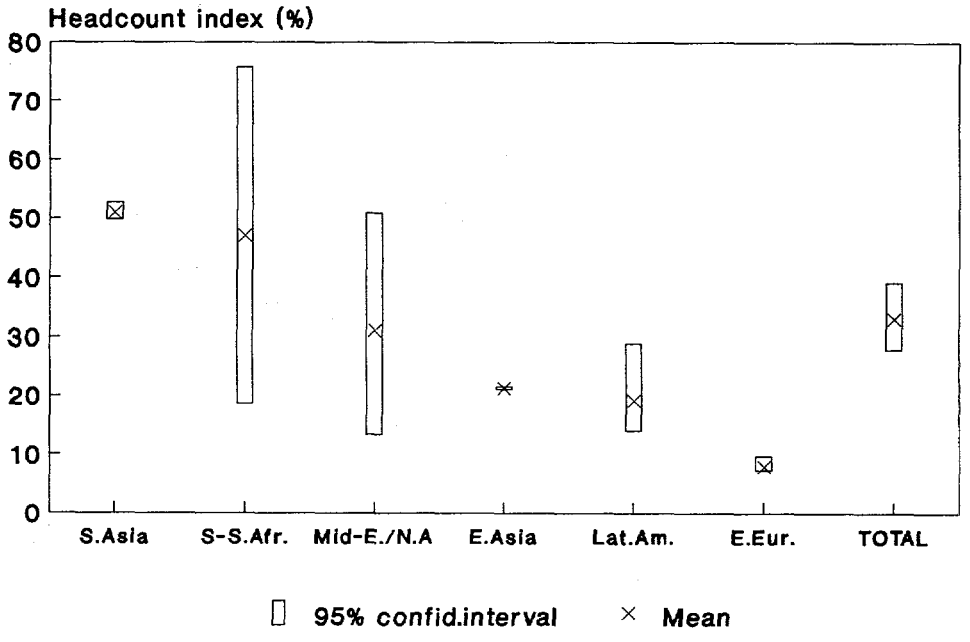

Note: Confidence intervals around point estimates of the percentage of the population of each region consuming less than $31 per month in 1985 at purchasing power parity. The confidence intervals allow only for the extrapolation errors when appropriate survey data were unavailable.

Source: Ravallion, Datt and van de Walle 1991.

The question often arises as to whether or not a developing country should bother to invest in a household survey. In principle, one way to answer this is by assessing the costs and benefits of such data, but in practice, although costs are easily assessed, benefits are not. There are qualitative gains in terms of the ability of the government, the governed, and the wider research community to critically yet systematically assess actual or proposed policies, and the country's overall progress. Without a poverty profile one can only guess at what pro-poor policies look like. What value does one attach to these gains, or, indeed, to any fundamental knowledge about the living standards of a society? I have never seen an attempt to do a cost-benefit analysis of

the case for collecting national accounts data; indeed, it is almost universally assumed that such data are as essential as government itself. Such an assumption would seem no less compelling when applied to basic national household surveys of consumption or income done periodically.

Armed with a comparable household consumption survey for the two or more dates or places being compared (or with and without some policy change), the best first step for the analyst is to construct each of the "poverty incidence curves", given by the proportion of the population estimated to be consuming less than a given consumption per person, while allowing the latter to vary from zero to the highest level. If the plotted poverty incidence curves do not cross each other at any interior point - a situation called "first-order dominance" - then the qualitative poverty comparison is unambiguous, and further calculations are quite unnecessary. No matter where the poverty line is drawn or which poverty measure is used (within a very broad class of measures), the situation with the lower poverty incidence curve will have less poverty.

If the two curves do cross each other, or one requires a quantitative poverty comparison, then more information is needed. One can add information about what is considered to be a reasonable poverty line in the country, and/or one can add information about what are considered to be reasonable properties of the poverty measure. Both sorts of information can be contentious.

I have surveyed the main methods used for determining poverty lines when it is necessary to do so. One of the most common, and quite defensible, methods is to use the cost of a bundle of goods deemed to be adequate to guarantee that the nutritional and other basic needs of a typical person are met. In making comparisons of absolute poverty one should only adjust the poverty line so that it maintains constant value in terms of the living standards indicator being used; using consumption

or income, this typically means that it should only be adjusted for differences over time or space in the cost-of-living facing the poor.

However, in many instances one will find that poverty comparisons are relatively insensitive to the precise poverty line used over a reasonably wide range; while the precise poverty numbers will change (and possibly by quite a lot), the rankings need not. Provided one can say that the poverty line does not exceed some critical value, and that value does not occur above the lowest of the interior intersection points of the poverty incidence curves, then that is all one needs; again the qualitative comparison is conclusive whatever the poverty measure.

But, if one does not know enough about the poverty line, then restrictions on the class of poverty measures will be needed in situations where first-order dominance does not hold, such as when some of the poor gain while others lose. It may well be enough to restrict attention to measures that have the property that an increase in any person's consumption cannot increase poverty. That will be sufficient for deciding the qualitative poverty comparison if the "poverty deficit curves", given by the area under the poverty incidence curves, do not cross each other at any interior point. In other words, one of them is higher than the other for at least some points and is nowhere lower up to the maximum poverty line. Then second-order dominance holds. For example, in a comparison in which some of the poor lose, while others gain, if aggregate gains exceed losses then the poverty deficit curve will have fallen.

If that test is not conclusive (the poverty deficit curves intersect at some point below the highest admissible poverty line), then one will need to say something about how the poverty measure weights gains or losses at different consumption levels amongst the poor, and, in particular, how sensitive it is the severity of poverty. A natural assumption to make is that a higher weight is attached to given gain in

consumption or income the poorer the gainer is. This requires a comparison of the "poverty severity curves", given by the areas under the poverty deficit curves, and it yields a "third-order dominance" test, analogous to the first and second order tests described above. If it is passed, then this implies an unambiguous poverty comparison for measures of the severity of poverty.

A purely qualitative poverty comparison will not always be enough. We may want to know how much the living standards of the poor have changed, and what factors contributed to that change. There are decomposition methods which can help here. A poverty profile can identify the sub-groups where poverty incidence, depth, or severity are highest, for the purposes of better understanding how actual or expected changes will affect the poor. Comparisons over time can also be decomposed into various components, such as due to growth versus redistribution, or sectoral components. And, at the finest level of decomposition, likely welfare impacts at the household level can often be quantified with reasonable confidence using the tools of modern micro-econometrics. This can give great scope for a deeper understanding of impacts on the poor, and their likely behavioral responses.

I have given a number of illustrations of how these methods can be used in routine poverty assessments and policy analysis. It has been seen that some assessments of impacts on poverty are far less robust than others to the assumptions made in measuring poverty. Similarly, some policy comparisons rest more heavily on the often difficult value judgements made in measuring poverty than do others. There is nothing general one can say on this.

What is clear is that a few relatively simple tools for poverty analysis can greatly facilitate assessments of how robust poverty comparisons are to measurement assumptions. At least the worst outcome can be avoided, namely that of not knowing how fragile policy

conclusions really are to the data and assumptions on which they are based. But, as some of the applications studied here have illustrated, one may often find that the poverty comparison is robust over a wide range of circumstances, adding confidence to the policy conclusion.

Appendix: Some Formal Definitions for Continuous Distributions

First we define the three poverty measures discussed in section 2.5. The cumulative distribution function is defined by:

$$\text{(A1)} \qquad F(y) = \int_0^y f(x)dx$$

which is the probability of observing someone with a standard of living less than y (where f(x) is the probability of observing a living standards indicator with the value x.) Letting z denote the poverty line, the general class of FGT poverty measures can be written in the form:

$$\text{(A2)} \qquad P_\alpha(z) = \int_0^z (1-x/z)^\alpha \, f(x)dx$$

The head-count index H(z) is then obtained when α=0. When α=1, (A2) becomes the poverty-gap index PG(z), while when α=1, it is the distribution sensitive measure $P_2(z)$.

Next we define the three corresponding poverty curves described in section 2.7. Instead of the poverty line z being a single number, consider the widest range of possible poverty lines from 0 to the maximum possible level z^{max}. The "poverty incidence curve" is simply defined by the values taken by the cumulative distribution function over the relevant interval; thus the poverty incidence curve is F(z) as z varies from 0 to z^{max}. Points on the poverty incidence curve can be read off as the head-count index of poverty i.e., H(z) = F(z).

The "poverty deficit curve" is defined by:

$$\text{(A3)} \qquad D(z) = \int_0^z (z-x)f(x)dx = \int_0^z F(x)dx$$

as z varies from 0 to z^{max}. The poverty gap index at any poverty line can be obtained directly from the points on the poverty deficit curve using the fact that:[60]

$$PG(z) = D(z)/z$$

which also gives a simple formula for the mean living standard of those with less than z, μ^z, namely

$$\mu^z = z - D(z)/F(z)$$

Similarly, the mean living standard of the poorest p percent of the population μ_p is simply:

$$\mu_p = z^* - D(z^*)/p$$

where the point z^* is defined implicitly by $F(z^*)=p$.

Again letting z vary from 0 to z^{max}, the "poverty severity curve" is defined by:

$$S(z) = \int_0^z (z-x)F(x)dx = \int_0^z D(x)dx \qquad (A4)$$

Points on the poverty severity curve can also be used to derive the P_2 poverty measure using the fact that:

$$P_2(z) = 2S(z)/z^2$$

---

60 This formula and the one for P_2 following are obtained by integrating A1 by parts and using a little algebra.

It can be seen that each of these curves is obtained by integrating the preceding one over $(0, z^{max})$; see Figure 4. Thus, if the poverty deficit curve (for example) of one distribution lies entirely above another then that will also be true of the poverty severity curves. However, that will not necessarily be true of the poverty incidence curves. Thus an unambiguous ranking of distributions in terms of P_α poverty measures for all poverty lines up to z^{max} must also imply an unambiguous ranking for all higher values of α over the same range of poverty lines.

References

Akerlof, George A., An Economic Theorist's Book of Tales, Cambridge: Cambridge University Press, 1984.

Anand, Sudhir and Christopher Harris, "Food and Standard of Living: An Analysis on Sri Lankan Data". In Jean Dreze and Amartya Sen (eds), The Political Economy of Hunger, Volume 1, Entitlement and Well-Being, Oxford, Oxford University Press, 1991.

Arrow, Kenneth, Social Choice and Individual Values, 2nd edition, New York: Wiley, 1963.

Ashenfelter, Orley, Angus Deaton, and Gary Solon, "Collecting Panel Data in Developing Countries: Does it Make Sense?", Living Standards Measurement Study Working Paper No.23, World Bank, Washington DC, 1986.

Atkinson, Anthony B., "On the Measurement of Inequality", Journal of Economic Theory Vol.2, 1970, pp:244-263.

_____., The Economics of Inequality, Oxford: Oxford University Press, 1975.

_____, "On the Measurement of Poverty", Econometrica Vol. 55, 1987, pp: 749-64.

_____, "Measuring Poverty and Differences in Family Composition", London School of Economics, 1988.

_____, Poverty and Social Security, New York: Harvester Wheatsheaf, 1989.

_____, "Comparing Poverty Rates Internationally: Lessons from Recent Studies in Developed Countries", World Bank Economic Review Vol.5, 1991, pp: 3-22.

Atkinson, Anthony B., and F. Bourguignon, "The Comparison of Multi-Dimensional Distributions of Economic Status", Review of Economic Studies Vol. 49, 1982, pp: 183-201.

_____ and _____, "Income Distribution and Differences in Needs", in George R. Feiwel (ed) Arrow and the Foundations of the Theory of Economic Policy, London: Macmillan Press, 1987.

_____ and _____, "The Design of Direct Taxation and Family Benefits", Journal of Public Economics Vol. 41, 1989, pp: 3-29.

Beaton, George H., "Energy in Human Nutrition: Perspectives and Problems", Nutrition Reviews Vol 41, 1983, pp: 325-340.

Besley, Timothy, and Steven Coate, "Workfare vs. Welfare: Incentive Arguments for Work Requirements in Poverty Alleviation Programs", Discussion Paper 142, Research Program in Development Studies, Woodrow Wilson School of Public and International Affairs, Princeton University, 1988.

Bhargava, Alok, and Martin Ravallion, "Is Household Consumption a Martingale? Tests for Rural South India", mimeo, Welfare and Human Resources Division, World Bank, 1991.

Blundell, Richard, and Ian Preston, "The Distinction Between Income and Consumption in Measuring the Distribution of Household Welfare", mimeo, University College London, 1991.

Boateng, E. Oti, Kodwo Ewusi, Ravi Kanbur, and Andrew McKay, "A Poverty Profile for Ghana 1987-88", SDA Working Paper No. 5, World Bank, Washington DC.

Bouis, Howarth E., and Lawrence Haddad, Effects of Agricultural Commercialization on Land Tenure, Household Resource Allocation and Nutrition in the Philippines, Research Report 79, International Food Policy Research Institute, Washington DC.

Bourguignon, Francois, "Family Size and Social Utility. Income Distribution Dominance Criteria", Journal of Econometrics Vol. 42, 1989, pp: 67-80.

Box, George E.P., William G. Hunter, and J. Stuart Hunter, Statistics for Experimenters. An Introduction to Design, Data Analysis, and Model Building, New York: John Wiley and Sons, 1978.

Boyce, James, and Martin Ravallion, "A Dynamic Econometric Model of Agricultural Wage Determination in Bangladesh", Oxford Bulletin of Economics and Statistics Vol. 53, 1991, pp: 361-376.

Chaudhuri, Shubham, and Martin Ravallion, "How Well do Static Welfare Indicators Identify the Chronically Poor?", Living Standards Measurement Study Working Paper, in press.

Chen, Shaohua, Gaurav Datt, and Martin Ravallion, "POVCAL: A Program for Calculating Poverty Measures from Grouped Data", Welfare and Human Resources Division, World Bank.

Chow, K. Victor, Paul D. Thistle, and S. Chakraboti, "Testing for Stochastic Dominance: A General Approach", Paper presented at Eastern Economics Association Meeting, Pittsburgh, 1991.

Cowell, Frank A., Measuring Inequality. Oxford: Philip Allan, 1977.

Daniel, Wayne W., "Applied Nonparametric Statistics", Boston: PWS-KENT Publishing Company, 1990.

Datt, Gaurav, "Computational Tools for Poverty Measurement and Analysis Using Grouped Data", Welfare and Human Resources Division, World Bank, 1991.

Datt, Gaurav and Martin Ravallion, "Regional Disparities, Targeting, and Poverty in India", in Michael Lipton and Jacques van der Gaag (eds), Including the Poor, World Bank, 1991a.

_____ and _____, "Growth and Redistribution Components of Changes in Poverty Measures: A Decomposition with Applications to Brazil and India in the 1980s", Living Standards Measurement Study Working Paper No.83, World Bank, Washington DC, 1991b.

_____ and _____, "Impacts on Time Allocation and Poverty of Public Works in Two Maharashtran Villages", Welfare and Human Resources Division, World Bank, Washington DC., 1991c.

Deaton, Angus, "The Measurement of Welfare. Theory and Practical Guide-Lines", Living Standards Measurement Study Working Paper No.7, World Bank, 1980.

_____, "Quality, Quantity and Spatial Variation of Price", American Economic Review Vol. 78, 1988, pp:418-430.

_____, "Rice Prices and Income Distribution in Thailand: A Non-Parametric Analysis", The Economic Journal Vol. 99 (Conference Supplement), 1989, pp:1-37.

_____, "Household Saving in LDC's: Credit Markets, Insurance, and Welfare", Discussion Paper 153, Research Program in Development Studies, Woodrow Wilson School of Public and International Affairs, Princeton University, 1991.

Deaton, Angus, and Dwayne Benjamin, "The Living Standards Survey and Price Policy Reform. A Study of Cocoa and Coffee Production in Cote d'Ivoire", Living Standards Measurement Study Working Paper No. 44, World Bank, Washington DC, 1988.

Deaton, Angus, and John Muellbauer, Economics and Consumer Behavior, Cambridge: Cambridge University Press, 1980.

_____, and _____, "On Measuring Child Costs: With Applications to Poor Countries", Journal of Political Economy Vol. 94, 1986, pp:720-744.

Dreze, Jean, and Amartya Sen, Hunger and Public Action, Oxford: Oxford University Press, 1989.

Fisher, Franklin M., "Household Equivalence Scales and Interpersonal Comparisons", Review of Economic Studies Vol. 54, 1987, pp: 519 -524.

Foster, James, "On Economic Poverty: A Survey of Aggregate Measures", Advances in Econometrics Vol.3, 1984, pp:215-251.

Foster, James., J. Greer, and E. Thorbecke, "A Class of Decomposable Poverty Measures", Econometrica Vol. 52, 1984, pp: 761-765.

Foster, James, and A. F. Shorrocks, "Poverty Orderings", Econometrica Vol. 56, 1988, pp: 173-177.

_____, and _____, "Subgroup Consistent Poverty Indices", Econometrica Vol. 59, 1991, pp:687-709.

Gaiha, Ragav, and Deolalikar, Anil B. "Permanent, Expected and Innate Poverty: Estimates for Semi-Arid Rural South India, 1975-84." Discussion Paper 89-17, Institute for Economic Research, University of Washington, Seattle, Washington, 1989.

Gautam, Madhur. "Measurement Errors and Consumption Analysis. A Case Study on Indian Panel Data". Mimeo, Welfare and Human Resources Division, World Bank, Washington DC, 1991.

Glewwe, Paul, "Improving Data on Poverty in the Third World: The World Bank's Living Standards Measurement Study", PRE Working Paper No. 416, Population and Human Resources Department, World Bank, Washington DC, 1990.

Glewwe, Paul and Jacques van der Gaag, "Identifying the Poor in Developing Countries: Do Different Definitions Matter?", World Development Vol. 18, 1990, pp: 803-814.

Greer, J., and Erik Thorbecke, "A Methodology for Measuring Food Poverty Applied to Kenya", Journal of Development Economics Vol. 24, 1986, pp: 59-74.

Grootaert, Christiaan, and Timothy Marchant, "The Social Dimensions of Adjustment Priority Survey. An Instrument for the Rapid Identification and Monitoring of Policy Target Groups", Social Dimensions of Adjustment in Sub-Saharan Africa, Working Paper 12, World Bank, Washington DC.

Grosh, Margaret E., "The Household Survey as a Tool for Policy Change", Living Standards Measurement Study Working Paper No. 80, World Bank, Washington DC.

Hadar, Josef and William R. Russell, "Rules for Ordering Uncertain Prospects", American Economic Review Vol. 59, 1969, pp:25-34.

Haddad, Lawrence and John Hoddintott, "Gender Aspects of Household Expenditures and Resource Allocation in the Cote d'Ivoire", mimeo, Centre for the Study of African Economies, Oxford University.

Haddad, Lawrence and Ravi Kanbur, "How Serious is the Neglect of Intra-Household Inequality?", The Economic Journal Vol. 100, 1990, pp: 866-881.

Hagenaars, Alsi, and Klaas de Vos, "The Definition and Measurement of Poverty", The Journal of Human Resources Vol 23, 1988, pp: 211-221.

Hagenaars, Aldi J.M., and Bernard M.S. van Praag, "A Synthesis of Poverty Line Definitions", Review of Income and Wealth, 1985, pp:139-154.

Hamburg, Morris, Statistical Analysis for Decision Making, New York: Harcourt Brace Jovanovoch, 1977.

Huppi, Monika, and Martin Ravallion, "The Sectoral Structure of Poverty During an Adjustment Period. Evidence for Indonesia in the Mid-1980s", World Development 1991.

Jorgenson, Dale W., and D.T. Slesnick, "Aggregate Consumer Behavior and the Measurement of Inequality", Review of Economic Studies Vol. 60, 1984, pp: 369-392.

Kakwani, Nanak, Income Inequality and Poverty. Methods of Estimation and Policy Applications, Oxford: Oxford University Press, 1980.

_____, "On a Class of Poverty Measures", Econometrica Vol 48, 1980, pp: 437-446.

_____, "On Measuring Undernutrition", Oxford Economic Papers Vol 41, 1989, pp:528-552.

_____, "Testing for the Significance of Poverty Differences With Application to Cote d'Ivoire", Living Standards Measurement Study Working Paper No. 62, World Bank, 1990.

Kanbur, Ravi, "Measurement and Alleviation of Poverty", IMF Staff Papers Vol.36, 1987, pp. 60-85.

_____, "Poverty and the Social Dimensions of Structural Adjustment in Cote d'Ivoire", Social Dimensions of Adjustment in Sub-Saharan Africa, Policy Analysis, World Bank, Washington DC.

Kapteyn, Arie, Peter Kooreman, and Rob Willemse, "Some Methodological Issues in the Implementation of Subjective Poverty Definitions", The Journal of Human Resources Vol 23, 1988, pp 222-242.

King, Mervyn A., "Welfare Analysis of Tax Reforms Using Household Level Data", Journal of Public Economics Vol 21, 1983, pp: 183-214.

Lanjouw, Peter, and Nicholas Stern, "Poverty in Palanpur", World Bank Economic Review Vol. 5, 1991, pp: 23-56.

Levy, Paul S., and Stanley Lemeshow, Sampling of Populations. Methods and Applications, New York: John Wiley and Sons, 1991.

Lipton, Michael, Why Poor People Stay Poor, London: Temple Smith, 1977.

_____, "Poverty, Undernutrition, and Hunger". World Bank Staff Working Paper 597, 1983, World Bank, Washington DC.

_____, "The Poor and the Poorest. Some Interim Findings". World Bank Discussion Paper 25, 1988, World Bank, Washington DC.

Lipton, Michael, and Jacques van der Gaag, "Poverty: A Research and Policy Framework", in Michael Lipton and Jacques van der Gaag (eds) Including the Poor, 1991.

McElroy, Marjorie B., "The Empirical Content of Nash-Bargained Household Behavior", The Journal of Human Resources Vol. 25, 1990, pp:559-583.

Meerman, Jacob, Public Expenditure in Malaysia. Who Benefits and Why. Oxford: Oxford University Press, 1979.

Muellbauer, John, "Inequality Measures, Prices and Household Composition", Review of Economic Studies Vol. 41, 1974, pp: 493 -504.

Ng, K., "Welfarism: A Defence Against Sen's Attack", The Economic Journal Vol.91, 1981.

Orshansky, Mollie, "Counting the Poor: Another Look at the Poverty Profile", Social Security Bulletin Vol.28, 1965, pp: 3-29.

Pashardes, Panos, "Contemporaneous and Intertemporal Child Costs. Equivalent Expenditure vs. Equivalent Income Scales", Journal of Public Economics Vol. 45, 1991, pp. 191-213.

Pollak, R.A., and T.J. Wales, "Welfare Comparisons and Equivalence Scales", American Economic Review Vol.69, 1979, pp:216-221.

Quirk, J., and R. Saposnik, "Admissibility and Measurable Utility Functions", Review of Economic Studies Vol.29, 1962, pp:140-146.

Ravallion, Martin, "Expected Poverty under Risk-Induced Welfare Variability", The Economic Journal Vol. 98, 1988, pp: 1171-1182.

_____, "Land-Contingent Poverty Alleviation Schemes", World Development Vol. 17, 1989, pp: 1223-1233.

_____, "The Challenging Arithmetic of Poverty in Bangladesh", The Bangladesh Development Studies Vol 43, 1990a, pp: 35-53.

_____, "Rural Welfare Effects of Food Price Changes Under Induced Wage Responses: Theory and Evidence for Bangladesh", Oxford Economic Papers Vol. 42, 1990b, pp 574-585.

_____, "Reaching the Rural Poor through Public Employment: Arguments, Experience and Lessons from South Asia", World Bank Research Observer 1991a.

_____, "On the Coverage of Public Employment Schemes for Poverty Alleviation", Journal of Development Economics Vol.34, pp. 57-79, 1991b.

_____, "Does Undernutrition Respond to Incomes and Prices: Dominance Tests for Indonesia", World Bank Economic Review Vol. 6, 1992a.

_____, "Poverty Alleviation through Regional Targeting: A Case Study for Indonesia", in Avishay Braverman, Karla Hoff, and Joseph Stiglitz (eds) The Economics of Rural Organization, 1992b.

_____, "On Hunger and Public Action: A Review Article on a Book by Jean Dreze and Amartya Sen", World Bank Research Observer 1992c.

Ravallion, Martin and Shubham Chaudhuri, "Testing Risk-Sharing in Three Indian Villages", Welfare and Human Resources Division, World Bank, 1991.

Ravallion, Martin and Monika Huppi, "Measuring Changes in Poverty: A Methodological Case Study of Indonesia During an Adjustment Period", World Bank Economic Review Vol. 5, 1991, pp:57-84.

Ravallion, Martin and Dominique van de Walle, "The Impact on Poverty of Food Pricing Reforms: A Welfare Analysis for Indonesia" Journal of Policy Modeling Vol. 13, 1991a, pp: 281-299.

_____ and _____, "Urban-Rural Cost-of-Living Differentials in a Developing Economy", Journal of Urban Economics Vol. 29, 1991b, pp: 113-127.

Ravallion, Martin, Gaurav Datt, and Dominique van de Walle, "Quantifying Absolute Poverty in the Developing World", Review of Income and Wealth 1991, in press.

Rodgers, John L., and Joan R. Rodgers, "Measuring the Intensity of Poverty Among Subpopulations. Applications to the United States", The Journal of Human Resources Vol. 26, 1991, pp: 338-361.

Rothschild, Michael, and Joseph E. Stiglitz. "Increasing Risk: 1. A Definition", Journal of Economic Theory Vol. 2, 1970, pp: 225-243.

Schultz, T. Paul., "Testing the Neoclassical Model of Family Labor Supply and Fertility", The Journal of Human Resources Vol. 25, 1990, pp:599-634.

Sen, Amartya, "Poverty: An Ordinal Approach to Measurement", Econometrica Vol. 46, 1976, pp:437-446.

_____, "Personal Utilities and Public Judgements: Or What's Wrong with Welfare Economics?", The Economic Journal Vol. 89, 1979, pp: 537-558.

_____, Poverty and Famines: An Essay on Entitlement and Deprivation, Oxford: Oxford University Press, 1981.

_____, Resources, Values and Development, Oxford: Basil Blackwell, 1984.

_____, The Standard of Living, Cambridge: Cambridge University Press, 1987.

_____. "Development as Capability Expansion", in Human Development in the 1980s and Beyond. Journal of Development Planning, Department of International Economic and Social Affairs, United Nations, No.19, 1989, pp:41-58.

Shorrocks, Anthony F. "Ranking Income Distributions", Economica Vol.50, 1983, pp: 3-17.

Slesnick, Daniel T. "Consumption and Poverty", Processed, Kennedy School of Government, Harvard University, 1990.

Thomas, Duncan, "Intrahousehold Resource Allocation: An Inferential Approach", Journal of Human Resources 1990.

Tobin, James, "On Limiting the Domain of Inequality", Journal of Law and Economics Vol 13, 1970, pp: 263-277.

Townsend, Robert, "Risk and Insurance in Village India". Department of Economics, University of Chicago, 1990.

United Nations, National Household Survey Capability Programme. Household Income and Expenditure Surveys: A Technical Study. UN Department of Technical Co-Operation for Development and Statistical Office, New York, 1989.

United Nations Development Programme (UNDP), Human Development Report, Oxford: Oxford University Press, 1990.

van de Walle, Dominique, "Poverty and Inequality in Latin America and the Caribbean During the 1970s and 1980s: An Overview of the Evidence", A View from LATHR, No. 22, Human Resources Division, Technical Department, Latin America and the Caribbean Region, The World Bank, 1991.

van der Gaag, Jacques, Elene Maknonnen, and Pierre Englebert, "Trends in Social Indicators and Social Sector Spending", PRE Working Paper 662, World Bank, Washington DC, 1991.

Varian, Hal R., Microeconomic Analysis (2nd edition), New York: W.W. Norton.

Vartia, Y.O., "Efficient Methods of Measuring Welfare Changes and Compensated Income in Terms of Ordinary Demand Functions", Econometrica Vol. 51, 1983, pp: 79-98.

Villasenor, J., and B. C. Arnold, "Elliptical Lorenz Curves", Journal of Econometrics Vol. 40, 1989, pp: 327-338.

World Bank, The World Bank Annual Report 1990, Washington DC: The World Bank, 1990a.

_____, World Development Report 1990. Poverty, Oxford: Oxford University Press, 1990b.

_____, Indonesia. Strategy for a Sustained Reduction in Poverty, World Bank, Washington DC, 1991.

World Health Organization, Energy and Protein Requirements, WHO Technical Report Series 724, Geneva, 1985.

Distributors of World Bank Publications

ARGENTINA
Carlos Hirsch, SRL
Galeria Guemes
Florida 165, 4th Floor-Ofc. 453/465
1333 Buenos Aires

AUSTRALIA, PAPUA NEW GUINEA, FIJI, SOLOMON ISLANDS, VANUATU, AND WESTERN SAMOA
D.A. Books & Journals
648 Whitehorse Road
Mitcham 3132
Victoria

AUSTRIA
Gerold and Co.
Graben 31
A-1011 Wien

BAHRAIN
Bahrain Research and Consultancy Associates Ltd.
P.O. Box 22103
Manama Town 317

BANGLADESH
Micro Industries Development Assistance Society (MIDAS)
House 5, Road 16
Dhanmondi R/Area
Dhaka 1209

Branch offices:
Main Road
Masjdee Court
Noakhali - 3800

76, K.D.A. Avenue
Kulna

BELGIUM
Jean De Lannoy
Av. du Roi 202
1060 Brussels

CANADA
Le Diffuseur
C.P. 85, 1501B rue Ampère
Boucherville, Québec
J4B 5E6

CHINA
China Financial & Economic Publishing House
8, Da Fo Si Dong Jie
Beijing

COLOMBIA
Infoenlace Ltda.
Apartado Aereo 34270
Bogota D.E.

COTE D'IVOIRE
Centre d'Edition et de Diffusion Africaines (CEDA)
04 B.P. 541
Abidjan 04 Plateau

CYPRUS
MEMRB Information Services
P.O. Box 2098
Nicosia

DENMARK
SamfundsLitteratur
Rosenoerns Allé 11
DK-1970 Frederiksberg C

DOMINICAN REPUBLIC
Editora Taller, C. por A.
Restauración e Isabel la Católica 309
Apartado Postal 2190
Santo Domingo

EL SALVADOR
Fusades
Avenida Manuel Enrique Araujo #3530
Edificio SISA, 1er. Piso
San Salvador

EGYPT, ARAB REPUBLIC OF
Al Ahram
Al Galaa Street
Cairo

The Middle East Observer
41, Sherif Street
Cairo

FINLAND
Akateeminen Kirjakauppa
P.O. Box 128
SF-00101
Helsinki 10

FRANCE
World Bank Publications
66, avenue d'Iéna
75116 Paris

GERMANY
UNO-Verlag
Poppelsdorfer Allee 55
D-5300 Bonn 1

GREECE
KEME
24, Ippodamou Street Platia Plastiras
Athens-11635

GUATEMALA
Librerias Piedra Santa
5a. Calle 7-55
Zona 1
Guatemala City

HONG KONG, MACAO
Asia 2000 Ltd.
48-48 Wyndham Street
Winning Centre
2nd Floor
Central Hong Kong

INDIA
Allied Publishers Private Ltd.
751 Mount Road
Madras - 600 002

Branch offices:
15 J.N. Heredia Marg
Ballard Estate
Bombay - 400 038

13/14 Asaf Ali Road
New Delhi - 110 002

17 Chittaranjan Avenue
Calcutta - 700 072

Jayadeva Hostel Building
5th Main Road Gandhinagar
Bangalore - 560 009

3-5-1129 Kachiguda Cross Road
Hyderabad - 500 027

Prarthana Flats, 2nd Floor
Near Thakore Baug, Navrangpura
Ahmedabad - 380 009

Patiala House
16-A Ashok Marg
Lucknow - 226 001

INDONESIA
Pt. Indira Limited
Jl. Sam Ratulangi 37
P.O. Box 181
Jakarta Pusat

ITALY
Licosa Commissionaria Sansoni SPA
Via Benedetto Fortini, 120/10
Casella Postale 552
50125 Florence

JAPAN
Eastern Book Service
37-3, Hongo 3-Chome, Bunkyo-ku 113
Tokyo

KENYA
Africa Book Service (E.A.) Ltd.
P.O. Box 45245
Nairobi

KOREA, REPUBLIC OF
Pan Korea Book Corporation
P.O. Box 101, Kwangwhamun
Seoul

KUWAIT
MEMRB Information Services
P.O. Box 5465

MALAYSIA
University of Malaya Cooperative Bookshop, Limited
P.O. Box 1127, Jalan Pantai Baru
Kuala Lumpur

MEXICO
INFOTEC
Apartado Postal 22-860
14060 Tlalpan, Mexico D.F.

MOROCCO
Société d'Etudes Marketing Marocaine
12 rue Mozart, Bd. d'Anfa
Casablanca

NETHERLANDS
InOr-Publikaties b.v.
P.O. Box 14
7240 BA Lochem

NEW ZEALAND
Hills Library and Information Service
Private Bag
New Market
Auckland

NIGERIA
University Press Limited
Three Crowns Building Jericho
Private Mail Bag 5095
Ibadan

NORWAY
Narvesen Information Center
Book Department
P.O. Box 6125 Etterstad
N-0602 Oslo 6

OMAN
MEMRB Information Services
P.O. Box 1613, Seeb Airport
Muscat

PAKISTAN
Mirza Book Agency
65, Shahrah-e-Quaid-e-Azam
P.O. Box No. 729
Lahore 3

PERU
Editorial Desarrollo SA
Apartado 3824
Lima

PHILIPPINES
International Book Center
Fifth Floor, Filipinas Life Building
Ayala Avenue, Makati
Metro Manila

POLAND
ORPAN
Palac Kultury i Nauki
00-901 Warszawa

PORTUGAL
Livraria Portugal
Rua Do Carmo 70-74
1200 Lisbon

SAUDI ARABIA, QATAR
Jarir Book Store
P.O. Box 3196
Riyadh 11471

MEMRB Information Services
Branch offices:
Al Alsa Street
Al Dahna Center
First Floor
P.O. Box 7188
Riyadh

Haji Abdullah Alireza Building
King Khaled Street
P.O. Box 3969
Damman

33, Mohammed Hassan Awad Street
P.O. Box 5978
Jeddah

SINGAPORE, TAIWAN, MYANMAR, BRUNEI
Information Publications Private, Ltd.
02-06 1st Fl., Pei-Fu Industrial Bldg.
24 New Industrial Road
Singapore 1953

SOUTH AFRICA, BOTSWANA
For single titles:
Oxford University Press Southern Africa
P.O. Box 1141
Cape Town 8000

For subscription orders:
International Subscription Service
P.O. Box 41095
Craighall
Johannesburg 2024

SPAIN
Mundi-Prensa Libros, S.A.
Castello 37
28001 Madrid

Librería Internacional AEDOS
Consell de Cent, 391
08009 Barcelona

SRI LANKA AND THE MALDIVES
Lake House Bookshop
P.O. Box 244
100, Sir Chittampalam A. Gardiner Mawatha
Colombo 2

SWEDEN
For single titles:
Fritzes Fackbokforetaget
Regeringsgatan 12, Box 16356
S-103 27 Stockholm

For subscription orders:
Wennergren-Williams AB
Box 30004
S-104 25 Stockholm

SWITZERLAND
For single titles:
Librairie Payot
6, rue Grenus
Case postale 381
CH 1211 Geneva 11

For subscription orders:
Librairie Payot
Service des Abonnements
Case postale 3312
CH 1002 Lausanne

TANZANIA
Oxford University Press
P.O. Box 5299
Dar es Salaam

THAILAND
Central Department Store
306 Silom Road
Bangkok

TRINIDAD & TOBAGO, ANTIGUA BARBUDA, BARBADOS, DOMINICA, GRENADA, GUYANA, JAMAICA, MONTSERRAT, ST. KITTS & NEVIS, ST. LUCIA, ST. VINCENT & GRENADINES
Systematics Studies Unit
#9 Watts Street
Curepe
Trinidad, West Indies

UNITED ARAB EMIRATES
MEMRB Gulf Co.
P.O. Box 6097
Sharjah

UNITED KINGDOM
Microinfo Ltd.
P.O. Box 3
Alton, Hampshire GU34 2PG
England

VENEZUELA
Libreria del Este
Aptdo. 60.337
Caracas 1060-A

YUGOSLAVIA
Jugoslovenska Knjiga
P.O. Box 36
Trg Republike
YU-11000 Belgrade

LSMS Working Papers (continued)

No. 54 *The Willingness to Pay for Education in Developing Countries: Evidence from Rural Peru*

No. 55 *Rigidité des salaires: Données microéconomiques et macroéconomiques sur l'ajustement du marché du travail dans le secteur moderne* (in French only)

No. 56 *The Poor in Latin America during Adjustment: A Case Study of Peru*

No. 57 *The Substitutability of Public and Private Health Care for the Treatment of Children in Pakistan*

No. 58 *Identifying the Poor: Is "Headship" a Useful Concept?*

No. 59 *Labor Market Performance as a Determinant of Migration*

No. 60 *The Relative Effectiveness of Private and Public Schools: Evidence from Two Developing Countries*

No. 61 *Large Sample Distribution of Several Inequality Measures: With Application to Côte d'Ivoire*

No. 62 *Testing for Significance of Poverty Differences: With Application to Côte d'Ivoire*

No. 63 *Poverty and Economic Growth: With Application to Côte d'Ivoire*

No. 64 *Education and Earnings in Peru's Informal Nonfarm Family Enterprises*

No. 65 *Formal and Informal Sector Wage Determination in Urban Low-Income Neighborhoods in Pakistan*

No. 66 *Testing for Labor Market Duality: The Private Wage Sector in Côte d'Ivoire*

No. 67 *Does Education Pay in the Labor Market? The Labor Force Participation, Occupation, and Earnings of Peruvian Women*

No. 68 *The Composition and Distribution of Income in Côte d'Ivoire*

No. 69 *Price Elasticities from Survey Data: Extensions and Indonesian Results*

No. 70 *Efficient Allocation of Transfers to the Poor: The Problem of Unobserved Household Income*

No. 71 *Investigating the Determinants of Household Welfare in Côte d'Ivoire*

No. 72 *The Selectivity of Fertility and the Determinants of Human Capital Investments: Parametric and Semiparametric Estimates*

No. 73 *Shadow Wages and Peasant Family Labor Supply: An Econometric Application to the Peruvian Sierra*

No. 74 *The Action of Human Resources and Poverty on One Another: What We Have Yet to Learn*

No. 75 *The Distribution of Welfare in Ghana, 1987–88*

No. 76 *Schooling, Skills, and the Returns to Government Investment in Education: An Exploration Using Data from Ghana*

No. 77 *Workers' Benefits from Bolivia's Emergency Social Fund*

No. 78 *Dual Selection Criteria with Multiple Alternatives: Migration, Work Status, and Wages*

No. 79 *Gender Differences in Household Resource Allocations*

No. 80 *The Household Survey as a Tool for Policy Change: Lessons from the Jamaican Survey of Living Conditions*

No. 81 *Patterns of Aging in Thailand and Côte d'Ivoire*

No. 82 *Does Undernutrition Respond to Incomes and Prices? Dominance Tests for Indonesia*

No. 83 *Growth and Redistribution Components of Changes in Poverty Measure: A Decomposition with Applications to Brazil and India in the 1980s*

No. 84 *Measuring Income from Family Enterprises with Household Surveys*

No. 85 *Demand Analysis and Tax Reform in Pakistan*

No. 86 *Poverty and Inequality during Unorthodox Adjustment: The Case of Peru, 1985–90*

No. 87 *Family Productivity, Labor Supply, and Welfare in a Low-Income Country*

The World Bank

Headquarters
1818 H Street, N.W.
Washington, D.C. 20433, U.S.A.

Telephone: (202) 477-1234
Facsimile: (202) 477-6391
Telex: WUI 64145 WORLDBANK
RCA 248423 WORLDBK
Cable Address: INTBAFRAD
WASHINGTONDC

European Office
66, avenue d'Iéna
75116 Paris, France

Telephone: (1) 40.69.30.00
Facsimile: (1) 40.69.30.66
Telex: 640651

Tokyo Office
Kokusai Building
1-1, Marunouchi 3-chome
Chiyoda-ku, Tokyo 100, Japan

Telephone: (3) 3214-5001
Facsimile: (3) 3214-3657
Telex: 26838