25534

Vulnerability and Poverty Measurement Issues for Public Policy

Jean-Yves Duclos

December 2002

Vulnerability and Poverty Measurement Issues for Public Policy

Jean-Yves Duclos

December 2002

WORLD BANK INSTITUTE

Promoting knowledge and learning for a better world

Social
Protection

LABOR MARKETS PENSIONS SOCIAL ASSISTANCE
THE WORLD BANK

Abstract

This paper presents the main approaches and tools to measuring poverty and vulnerability. This involves among other things the important issues of the setting of poverty lines, of the statistical use of surveys, of the empirical measurement of living standards, and of making poverty comparisons. The computation of aggregate poverty indices and the use of poverty dominance testing are also described. Throughout, implications for the understanding of the effects of public policy are highlighted, in particular those related to the design of safety nets.

Table of Contents

Vulnerability and Poverty Measurement Issues for Public Policy

Jean-Yves Duclos, CIRPÉE, Département d'économique, Universite Laval, Québec, Canada, G1K 7P4[1]

I. Measuring Poverty

Approaches to measuring poverty can be divided into two groups. Following Ravallion (1994), we will describe them as the welfarist and the non-welfarist approaches. The first approach tends to concentrate in practice mainly on comparisons of "economic well-being," which (for simplicity) we will also call "standard of living." As we will see, this approach has strong links with traditional economic theory, and it is also widely used by economists in the operations and research work of organizations such as the World Bank, the International Monetary Fund, and ministries of finance and planning of both developed and developing countries. The second approach has historically been advocated mainly by social scientists other than economists and partly as a reaction to the first approach. This second approach has nevertheless also been recently and increasingly suggested by economists and non-economists alike as a multidimensional complement to the classical standard of living approach.

The Welfarist Approach

The welfarist approach is strongly anchored in classical microeconomics, where, in the language of economists, "welfare" or "utility" are generally key in accounting for the behavior and the well-being of individuals. Classical microeconomics usually postulates that individuals are rational and that they can be presumed to be the best judges of the sort of life and activities that maximize their utility and happiness. Given their initial endowments (including time, talent, land, and capital), individuals make production and consumption choices using their set of preferences over bundles of consumption and production activities, and taking into account the available production technology and the consumer and producer prices that prevail in the economy. Under these assumptions and constraints, a process of individual free choice will maximize the individuals' utility; under additional (and admittedly somewhat restrictive) assumptions, a society of individuals all acting independently under this freedom of choice process will also lead to an outcome known as Pareto-efficient, in that no one's utility could be further improved by government intervention without decreasing someone else's utility.

Underlying the welfarist approach to poverty is the premise that account should be taken of the information revealed by individual behavior when it comes to assessing poverty. This premise holds, in particular, that the assessment of well-being should be consistent with the ordering of preferences revealed by the free choices made by individuals. For instance, a person could be observed to be poor by the total consumption or income standard of a poverty analyst. That same person could nevertheless be able (i.e., have the working capacity) to be non-poor. This could be revealed by the observation of a deliberate and free choice on the part of the individual to work and consume little, when the capability to work and consume more nevertheless exists. By choosing to spend little (possibly for the benefit of greater leisure), the person reveals that he is happier than if he worked and spent more. Although he could be considered poor by the standard of a (non-welfarist) poverty analyst, a welfarist judgement would conclude that this person is not poor. As we will discuss later, this can have important implications for the design and the assessment of public policy.

A pure welfarist approach faces important practical problems. To be operational, pure welfarism requires the observation of sufficiently informative revealed preferences. These are rarely easily observed and deciphered, however. For instance, for someone to be declared poor or not, it is not enough to know that person's current characteristics and living standard status, but it must also be inferred from the person's actions that *he* judges his utility status to be above a utility poverty level, as provided by a set of characteristics and a reference living standard. Another, more fundamental, problem with the pure welfarist approach is the need to measure levels of utility or "psychic happiness." How are we to measure the actual pleasure derived from experiencing economic well-being? Moreover, it is certainly difficult to attempt to compare that level of utility across individuals, and it is well known that such a procedure poses controversial ethical problems. Preferences are heterogeneous, personal characteristics and needs are diverse, households differ in size and composition, and prices vary across time and space. Besides, it is not clear that we should accept as ethically significant the actual level of utility felt by individuals. Why should a difficult-to-satisfy rich person be judged less happy than an easily-contented poor person? That is, why should the "grumbling rich" be judged "poorer" than a "contented peasant" (see Sen 1983, p. 160)?

Hence, welfarist comparisons of poverty almost invariably use imperfect but observable proxies for utilities, such as income or consumption. These money-metric indicators are often adjusted for differences in needs, prices, and household sizes and compositions, but they clearly do remain imperfect indicators of utility and well-being. Indeed, economic theory tells us little about how to use consumption or income to make interpersonal comparisons of well-being. Besides, the consumption and income proxies are also rarely able to take full account of the role of public goods and non-market commodities, such as safety, liberty, peace, and health, in determining well-being. In principle, such commodities can be valued using reference or "shadow" prices. In practice, this is difficult to do.

Non-Welfarist Approaches

There are two major non-welfarist approaches: the basic needs approach and the capability approach.

Basic Needs and Functionings

The first approach focuses on multidimensional *outcomes* that can generally be observed and monitored. These outcomes are usually (explicitly or implicitly) linked with the concept of *functionings*, a concept developed in Amartya Sen's influential work.

> *Living may be seen as consisting of a set of interrelated "functionings," consisting of beings and doings. A person's achievement in this respect can be seen as the vector of his or her functionings. The relevant functionings can vary from such elementary things as being adequately nourished, being in good health, avoiding escapable morbidity and premature mortality, etc., to more complex achievements such as being happy, having self-respect, taking part in the life of the community, and so on. (Sen 1997, p. 39)*

In this view, functionings can be understood to be constitutive elements of well-being. Hence, the functioning approach would generally not attempt to measure well-being in the single space (or in terms) of utility or happiness. This is partly because utility or happiness can be thought of as a single and reductive aggregate of functionings, which are multidimensional in nature. The functioning approach thus focuses on actual multidimensional achievements of definite, specific outcomes, such as the enjoyment of a particular type of commodity consumption or the *exercise* of rights. Among other things, these usually feature health, literacy, appropriate nutrition, having a shelter, not being in shame, to be well-clothed, to travel, to be involved in the community, etc.

The functioning approach is closely linked with the well-known *basic needs* approach, and the two are often difficult to distinguish in their practical application. Functionings, however, are not synonymous with basic needs, at least as typically defined and used in the empirical literature. More precisely, basic needs can be understood as the means that are required for individuals to achieve some functionings. Hence, basic needs are usually defined in terms of *means* rather than *outcomes*, for instance, as living in the proximity of providers of health care services (but not necessarily being in good health), as the number of years of achieved schooling (not necessarily as participating in the life of the community), and so on. In other words,

> *Basic needs may be interpreted in terms of minimum specified quantities of such things as food, shelter, water and sanitation that are necessary to prevent ill health, undernourishment and the like. (Streeten et al. 1981)*

Unlike functionings, which can be commonly defined for all individuals, the specification of basic needs depends on the characteristics of individuals and of the societies in which they live. For instance, the basic commodities required for someone to be in good health and not to be undernourished will depend on the climate and on the physiological characteristics of individuals. Similarly, the clothes necessary for one not to feel ashamed will depend on the norms of the society in which he lives, and the means necessary to travel, on whether he is handicapped or not. Hence, although the attainment of basic needs is an important element in assessing whether someone has achieved some functionings, this assessment must also use information on the characteristics and the socio-economic environment of the individual. Human diversity is such that equality in the space of basic needs generally translates into inequality in the space of functionings.

Whether unidimensional or multidimensional in nature, most applications of both the welfarist and the non-welfarist approaches to poverty measurement do recognize the role of needs and of the socio-economic environment in achieving well-being. Streeten et al. (1981) and others have nevertheless argued that the basic needs approach is less abstract than the welfarist approach in recognizing that role. As mentioned above, it can also be seen as a useful practical and operational step toward appraising the achievement of the more abstract "functionings." Clearly, however, there are important degrees in the achievement of basic needs and functionings. For instance, what does it mean precisely to be "adequately nourished"? What degree of nutrition adequacy is relevant for poverty assessment? Should the basic needs for adequate nutrition functioning only allow for the simplest possible diet and for the highest nutritional efficiency? Such an attitude is generally said to be overly restrictive and overly physiological in nature, but it is hard to find objective evidence to agree on the proper alternative. In addition, how ought we to understand such functionings as the functioning of self-respect? The appropriate width and depth of the concept of basic needs and functionings is admittedly ambiguous, as there are degrees of functionings which make life enjoyable in addition to being purely sustainable or satisfactory.

Capability Approach
A second alternative to the welfarist approach is called the *capability* approach, also pioneered and advocated in the last two decades by the work of Sen. The capability approach is defined by the capacity to achieve functionings, as defined above. In Sen's words (1997), the capability to function

> *represents the various combinations of functionings (beings and doings) that the person can achieve. Capability is, thus, a set of vectors of functionings, reflecting the person's freedom to lead one type of life or another. (p. 40)*

What matters for the capability approach is the *ability* of an individual to function well in society; it is not the actual functionings actually attained by the person. Having the capability to achieve "basic" functionings is the source of freedom to live well, and is thereby sufficient in the capability approach for one not to be poor or deprived.

The capability approach thus distances itself from achievements of specific outcomes or functionings. In this, it imparts considerable value to freedom of choice: a person will not be judged poor even if he chooses not to achieve some functionings, as long as he would be able to attain them if he so chose. This distinction between outcomes and the capability to achieve the outcomes also recognizes the importance of preference diversity and individuality in determining functioning choices. It is, for instance, not everyone's wish to be well-clothed or to participate in society, even if the capability is present.

An interesting example of the distinction between fulfillment of basic needs, functioning achievement, and capability is given by Townsend's (1979, table 6.3) deprivation index. The deprivation index is built from answers to questions such as whether someone "has not had an afternoon or evening out for entertainment in the last two weeks," or "has not had a cooked breakfast most days of the week." It may be, however, that one chooses deliberately not to go out for entertainment (he prefers to watch television), or that he chooses not to have a cooked breakfast (because he does not have time to prepare it), although he does have the capacity to do both. That person therefore achieves the functioning of being entertained

without meeting the basic need of going out once a fortnight, and does have the capacity to achieve the functioning of having a good breakfast, although he chooses not to.

The difference between the capability and the functioning or basic needs approach is in fact somewhat analogous to the difference between the use of income and consumption as standard-of-living indicators. Income shows the capability to consume, and "consumption functioning" can be understood as the outcome of the exercise of that capability. There is consumption only if a person chooses to enact his capacity to consume out of a given income. In the basic needs and functioning approach, deprivation comes from a lack of *direct* consumption or functioning experience; in the capability approach, poverty arises from the lack of incomes and capabilities, which are *indirectly* related to the actual functionings achieved.

Although the capability set is multidimensional, it thus exhibits a parallel with the unidimensional income indicator, whose size determines the size of the "budget set":

> *Just as the so-called "budget set" in the commodity space represents a person's freedom to buy commodity bundles, the "capability set" in the functioning space reflects the person's freedom to choose from possible livings. (Sen 1997, p. 40)*

This illustrates further the fundamental distinction between the space of achievements, the extents of freedoms and capabilities, and the resources required to generate these freedoms and to attain these achievements.

Measurement Difficulties
But how are we to measure capability? Unless a person chooses to enact them in the form of functioning achievements, capabilities are not observed directly. Achievement of all basic functionings implies non-deprivation in the space of all capabilities; but a failure to achieve all basic functionings does not imply capability deprivation. This makes the monitoring of functioning and basic needs an imperfect tool for the assessment of capability deprivation. Besides, and as for basic needs, there are clearly degrees of capabilities, some basic and some broader.

Non-welfarist (capability and basic needs) approaches to poverty measurement also suffer from some comparability problems. This is because they typically generate multidimensional *qualitative* poverty criteria: their fulfillment typically takes a simple dichotomic yes/no form. For instance, how should we assess adequately the degree of poverty of someone who has the capacity to achieve two functionings out of three, but not the third? Is that person necessarily "better off" than someone who can achieve only one, or even none of them? Are all capabilities of equal importance when we assess well-being? Money-metric indicators (such as income and consumption) are quantitative and more amenable in principle to subtle cardinal comparisons of well-being.

The multidimensionality of the non-welfarist criteria also translates into greater implementation difficulties than for the usual proxy indicators of the welfarist approach. In the welfarist approach, the size of the multidimensional budget is ordinarily summarized by income or total consumption, which can be thought of as a unidimensional indicator of freedom. A similar transformation into a unidimensional indicator is more difficult with the

5

capability and basic needs approaches. One possible solution is to use "efficiency-income units reflecting command over capabilities rather than command over goods and services" (Sen 1983, p. 343). This, however, is difficult to do in practice, since command over many capabilities is hard to translate in terms of a single indicator, and since the "budget units" are hardly comparable across functionings such as well-nourishment, literacy, feeling self-respect, and taking part in the life of the community. Furthermore, although there are many different combinations of consumption and functionings that are compatible with a unidimensional money-metric poverty threshold, the welfarist approach will generally not impose multidimensional thresholds. For instance, the welfarist approach will usually not require for one not to be poor that both food and non-food expenditures be larger than their respective food and non-food poverty lines. As indicated above, this simplifies the identification of the poor and the analysis of poverty.

II. Vulnerability and Poverty

The World Development Report 2000-2001 (World Bank, 2001) (WDR) defines poverty as "unacceptable deprivation in well-being." As outlined above, deprivation in well-being can be understood to arise from the inability to fulfill basic functionings, namely, deprivation in the space of basic capabilities. It has been argued in recent years that broader capabilities should be added to the basic material and physical capabilities typically retained by the traditional poverty assessments. These include, in particular, the capability of feeling secure and "empowered." Deprivation of these capabilities stems, respectively, from risk and vulnerability, and from voicelessness and powerlessness. In particular, vulnerability is the result of not being able to fulfill the basic functioning of security and of feeling at peace when pondering the future.

Vulnerability and insecurity are dynamic in nature. They are concerned with the anticipation of possible changes between now and the future. Insecurity is exposure to the risk of future events, some of them damageable to the person's well-being. Vulnerability is the possibility of suffering a decline in well-being, in particular a drop below some minimum benchmark or poverty threshold. This decline is brought about by shocks against which protection is either costly or not possible. The shocks can affect individuals, e.g., through loss of employment, accident, or death. They can also strike whole communities, such as villages, regions, or particular socio-economic groups. Examples of this include natural disasters, changes in export prices, and climactic and environmental changes.

Alwang and Siegel (2000) describe vulnerability as starting with the notion of *risk*, which is characterised by a probability distribution over events. Risk is compounded with the *exposure* of households—which is determined by such things as their geographical location, their assets, and their consumption and production choices—to provide the *hazard* which the household faces, and to which the household responds *ex ante* and *ex post*. *Ex ante* responses take the form of individual and social risk-reduction activities. *Ex post* responses are coping responses that decrease the harshness of the *outcome* of the hazard, which is a stochastic process. Vulnerability, then

> is defined with respect to events that may occur in the future. It is forward looking. While we can measure damage ex post...these are only the static outcomes of a continuous process of risk, exposure and response.

6

Vulnerability is the continuous forward-looking state. (Alwang and Siegel 2000, p. 5)

The measurement of vulnerability is particularly important for monitoring the well-being of the poor. Because of the poor's already difficult circumstances, the outcomes of vulnerability are typically harsher on the poor. The poor are also often more vulnerable because of their location and their characteristics (namely, their exposure), including a lower level of assets to protect them, less access to insurance and to input and output markets (to insulate them against idiosyncratic shocks), less access to public protection, and a lower level of empowerment. There are circumstances, however, in which the poor are less exposed to risk. This can be the case for near-subsistence households living in remote areas who are little (if at all) affected by market and price shocks (for a discussion and a test of this, see for instance Glewwe and Hall 1998).

Measuring vulnerability is also relevant for the design of poverty alleviation policies. A fundamental ethical and public policy issue is whether help should be targeted to the short-term poor, the long-term poor, or the most vulnerable among the poor and the non-poor. Presumably, if short-term poverty is such that it threatens physical subsistence, then observed, actual poverty should be the main target of public policy. A proper investigation of this issue must also assess the dynamic links between poverty and vulnerability, and how detrimental risk and variability of living standards are at the individual and social levels.

At the individual level, whether transitory poverty deserves as much policy attention as chronic poverty depends on the ability of individuals to buffer the harshness of temporary poverty. It is sometimes reported that transitory poverty and the efforts to cushion its effects affect both the short-term and the long-term prospects of individuals. The poor's limited resources hinder them from seeking adequate protection from short-term shocks. These limitations can take the form of a limited access to insurance markets (Jalan and Ravallion 1997) and to difficulties of putting aside resources for the future. There is also often on the part of the poor a trade-off between short-term protection against vulnerability and long-term exit from poverty. For instance, short-term protection against short-term poverty leads to use of physical and natural resources, whose depletion can be injurious to protection against long-term deprivation. This trade-off is particularly serious when the outcome of vulnerability brings the poor into abject deprivation. Besides, even if vulnerability were not instrumental in understanding the short-term and long-term prospects of the poor, it could still be of importance in as much as it affects one of the constitutive elements of individual well-being (fulfilling the basic functioning of security).

At the social level, risk and variability clearly generate a sense of social insecurity. However, they can allow some of the chronically poor to escape poverty temporarily, and some of the usually non-poor to experience poverty temporarily. This process can be argued to be a socially efficient outcome since it spreads the risk of being poor across individuals—the experience of deprivation is made more equal, and presumably more equitable. Such a sharing of deprivation across time is then conceptually close to the experience of social mobility, which is routinely deemed to be socially desirable. Assume, for instance, that there is a 50 percent chance of anyone being poor in a population at a given point in time. In the absence of social mobility and variability in poverty status across individuals, it will always be the same 50 percent of the population who will be poor. If, however, the risk of being poor at any point in time is spread equally across all, poverty will be shared across time

7

equally among all individuals. This would in turn generate more vulnerability, but it would still appear to be a socially more desirable distribution of poverty.

III. Measuring Vulnerability

As discussed above, not only is vulnerability a determinant and foreteller of future poverty, but it can also be argued to be an important element of actual well-being and deprivation, since it prevents the attainment of the functioning of security.

Vulnerability—which is forward-looking—has commonly been estimated either through indirect or past indicators of variability in well-being (defined alternatively as variability in utility, income or consumption, functionings, and capabilities). The indirect indicators are usually socio-economic proxies for risk exposure, such as living in cyclone-prone areas, being exposed to diseases, living in remote areas and thus far from health-care facilities, not owning land and capital (and therefore relying on potentially uncertain wage income for subsistence), etc. Indicators of past vulnerability commonly focus on the past variability of money-metric indicators such as consumption or income.

An interesting avenue in measuring past vulnerability is to consider movements of individuals in and out of poverty over a period of time, in comparison to the effect on these individuals of "chronic" or long-term poverty. Permanent or chronic poverty can be measured as poverty below a certain standard of living determined by some average long-term trend in living standards. This can be alternatively interpreted as expected poverty, if we think of individuals as using their past trend in living standards to predict their future poverty status. Intertemporal poverty can be carried out by summing over a period of time the effects of periodic lapses into poverty. When both intertemporal and long-term poverty are measured using a poverty index that is averse to variability of well-being, such as any of the distribution-sensitive poverty indices (for instance, the average squared poverty gap, or the Watts or the "Gini" indices of poverty—see Section VIII), the difference between intertemporal and chronic poverty gives an indicator of the effect of the "transiency" and variability of poverty status, and thus indirectly of the effect of vulnerability and insecurity on well-being. Intertemporal poverty is then the sum of chronic and transient poverty.

The temporal pattern of vulnerability and variability in well-being is also important. For instance, whether the temporal shocks are auto-correlated across time clearly matters for well-being. There may be no time to recuperate if shocks follow each other too quickly; this may lead to physiological and environmental fragility and to an inability to protect oneself against subsequent shocks, leading, for instance, to famines as opposed to transient short-term deprivation. Another important element is the correlation of the shocks across individuals and households. Are whole communities typically affected (in the case of epidemics, for instance), or are the shocks concentrated on just a few households or small socio-economic groups? The answers to these questions are important ingredients in the design of public policy, safety nets, and targeting schemes.

IV. Poverty Lines

Two major issues arise in the discussion of poverty lines. First, we must define the space in which poverty is measured. As outlined above, this can be the space of utility, basic needs,

functionings, or capabilities. Second, we must determine whether we are interested in an absolute or in a relative poverty line in the space considered.

Absolute and Relative Poverty Line

An absolute poverty line can be interpreted as fixed in any of the above spaces. A relative poverty line depends on the distribution of the living conditions (including the utilities, functionings, and capabilities) found in a society. Considerable controversy exists over whether absoluteness or relativity is a better property for a poverty threshold. Most analysts would nevertheless agree that a poverty threshold defined in the space of functionings and capabilities should be absolute (but even on this there is no unanimity). An absolute threshold in these spaces would, however, imply relativity of the corresponding thresholds in the space of the commodities and in the level of basic needs required to achieve these functionings.

There are two main reasons for this. First, the relative prices and the availability of commodities depend on the distribution of living standards. For instance, as a society initially develops, the affordability and accessibility of public transportation usually first increases as rising numbers of people need to travel to work and to trade, without first being able to afford the costs of private transportation. As societies become richer on average, however, their citizens make increasing use of private forms of transportation, which causes a drop in the supply and availability of public transportation, and to an increase in its price. This makes the capacity to travel (arguably an important functioning) more or less costly, depending on the state of economic development.

Second, not to be deprived of some capability may require the absence of relative deprivation in the space of some commodities. In support of this, there is Adam Smith's famous statement that the commodities needed to go without shame (an oft-mentioned basic functioning) can be to some extent relative to the distribution of living standards in a society.

> *By necessities, I understand not only the commodities which are indispensably necessary for the support of life but whatever the custom of the country renders it indecent for creditable people, even of the lowest order, to be without. (Smith 1776)*

Sen (1983) reinforces this by distinguishing clearly the two dimensions of capabilities and commodities.

> *I would like to say that poverty is an absolute notion in the space of capabilities but very often it will take a relative form in the space of commodities and characteristics. (p. 335)*

This has led some writers (particularly in developed countries) to conclude that attempts to preserve some degree of absoluteness in the space of commodities are untenable.

> *In summary, it does not seem possible to develop an approach to poverty measurement which is linked to absolute standards. While some analysts are uneasy with relativist concepts of poverty on the grounds that they are difficult to comprehend and can be seen as somewhat arbitrary and open to*

manipulation, no real practical alternative to relativist concepts exists.
(Saunders 1994, p. 227)

Social Exclusion and Relative Deprivation

Complete relativity of the poverty line in the space of commodities would nevertheless draw poverty analysis very close to the analysis of social exclusion (as exemplified by Rodgers et al. (1995) at the International Labour Organisation) and relative deprivation (as propounded for instance by Townsend 1979). Social exclusion entails "the drawing of inappropriate group distinctions between free and equal individuals which deny access to or participation in exchange or interaction" (Silver 1994, p. 557). This includes participation in property, earnings, public goods, and in the prevailing consumption level (Silver 1994, p. 541). Relative deprivation focuses on the inability to enjoy living standards and activities that are ordinarily observed in a society. Townsend (1979, p. 30) defines it as a situation in which

> *Individuals, families and groups in the population...lack the resources to obtain the types of diet, participate in the activities and have the living conditions and amenities which are customary or at least widely encouraged or approved, in the society to which they belong.*

Confounding absolute deprivation in the space of capabilities with relative deprivation in the space of commodities can, however, be a source of confusion in poverty comparisons. First, it tends to blur the operational and conceptual distinction between poverty and inequality. Second, it can hinder the identification of "core" or absolute poverty in any of the spaces. Identifying core poverty is, however, probably the most relevant task for discussions of public policy in developing countries. Monitoring and comparing poverty is a major purpose of poverty analysis, and this is made more difficult when a relative notion of poverty is employed. Third, although the moral appeal of Sen's capability approach has variously been invoked to justify the use of an entirely relative poverty line in the space of commodities, Sen himself does not accept this.

> *Indeed, there is an irreducible core of absolute deprivation in our idea of poverty, which translates reports of starvation, malnutrition and visible hardship into a diagnosis of poverty without having to ascertain first the relative picture. Thus the approach of relative deprivation supplements rather than supplants the analysis of poverty in terms of absolute dispossession. (Sen 1981, p.17)*

Furthermore,

> *...considerations of relative deprivation are relevant in specifying the "basic" needs, but attempts to make relative deprivation the sole basis of such specification is doomed to failure since there is an irreducible core of absolute deprivation in the concept of poverty. (Sen 1981, p. 17)*

Given the measurement difficulties involved in defining relative poverty lines that are adequately related to absolute poverty lines in the space of functionings and capabilities, economists often find the space of living standards to be the least problematic and clearest to use to define an absolute threshold below which individuals are considered poor. If this is

10

done, however, it must subsequently be admitted that the procedure will imply a set of thresholds in the space of functionings and capabilities that depends at least partly on the conditions of the society in which an individual lives. Indeed, for a given absolute level of living standard in the space of commodities, an individual's capabilities are relative, that is, they depend on the social environment, for functionings such as shamelessness and participation in the life of the community.

Estimating Poverty Lines

The methodology for estimating poverty lines has been most developed for identifying the resources required to fulfill basic physiological functionings. Although this methodology has often been set in a welfarist framework, it can equally be seen as an important tool for the basic needs, functioning, or capability approaches. The methodology has recently been most often applied in the context of developing countries.

Cost of Basic Food Needs

The procedure usually involves two steps. First, an estimation is made of the food expenditures that are necessary to achieve some minimally required level of food-energy or nutrient intake. These expenditures can be adjusted for differences in climate, region, or age. They can also be adjusted for variations in activity levels, although activity levels depend on the level of one's well-being, and thus on one's poverty status. Activity-level adjustments would thus involve a poverty line that evolves endogenously with the standard of living of individuals, a slightly awkward feature for comparing well-being across individuals and across time. This first step provides an estimate of a food poverty line. Early examples of the application of this approach include Rowntree (1901) and Orshansky (1965).

An important issue in this first step is whether variations in culinary tastes and food habits should be taken into account. If no account of such variations are taken, then a diet could be devised that minimizes food cost while providing the minimum required level of food-energy or nutrient intake. This would typically generate an unreasonably low level of expenditures, with an implied dietary basket of food commodities that could be far from those actually consumed by the poor.

If, however, full account of diversity in culinary tastes were to be taken, a serious risk would exist of overestimating the poverty lines of those individuals and groups of individuals with a greater taste for expensive foods (e.g., of high quality). This is commonly the case, for instance, for urban households, who customarily have more sophisticated culinary tastes than rural dwellers, and also greater access to a larger variety of imported and expensive foods. The procedure would then assign greater poverty lines to the urban versus the rural individuals. It would also mean that the equivalents of the food poverty lines in terms of living standards and "utilities" would depend on the peculiarities of the individuals' food preferences. This would clearly lead to inconsistent comparisons of well-being across urban and rural inhabitants, and would exaggerate the degree of poverty in urban as compared to rural areas.

Non-Food Poverty Lines

The subsequent step is to estimate a global poverty line (food and non-food). This is because it is normally admitted that an encompassing poverty analysis should cover both food and non-food expenditures. Standard household surveys do not, however, collect information on quantities for non-food items, and as a result it is difficult to estimate directly a non-food poverty line. The usual procedure for obtaining a global poverty line is to "scale up" the food

11

poverty line. The most popular method proceeds by dividing the food poverty line by the food share of total expenditures. The problem still remains of which food share to use. Popular practices make use of the average food share of a bottom proportion of the population (e.g., the poorest 25 percent), of those whose total expenditures equal the food poverty line, or of those whose food expenditures equal the food poverty line. Another method is to add to the food poverty line the expected non-food expenditures of those whose total expenditures equal the food poverty line. The choice among such adjustment procedures is inherently arbitrary, and the resulting estimate of the total poverty line will also be somewhat arbitrary. Moreover, some of the food shares will also vary with the distribution of living standards, especially when they are estimated over a range of individuals (such as the poorest 25 percent). To avoid inconsistencies in the comparisons of poverty, it is thus often suggested to use the same method and the same food share across the distributions being compared.

Food Energy Intake

A slightly different method for estimating poverty lines that is popular in the literature is the so-called Food-Energy-Intake (FEI) method. Estimates of the observed calorie intake of persons are first computed and then graphed against their observed (total or food) expenditures. The analyst then estimates the expenditures of those whose calorie intake is just at the minimum required for healthy subsistence. When these expenditures are on food, this provides a food poverty line, which can then be used as described above to provide an estimate of a global poverty line. When the expenditures are total expenditures, the FEI method provides a direct link between a minimum calorie intake and a global poverty line.

Relative and Subjective Poverty Lines

There are two other popular methodologies for the estimation of poverty lines. The first deals with relative poverty lines, which, as we saw above, can be useful to determine the commodities needed for "living without shame" and for participating in the "prevailing consumption level." A relative poverty line is typically set as an arbitrary proportion (often around 50 percent) of the median or the mean of living standards. Clearly, such a poverty line will vary with the central tendency of the distribution of living standards, and will not be the same across regions and time. One awkward feature of using a relative poverty line approach is that a policy which raises the living standards of all, but proportionately more those of the rich, will increase poverty, although the absolute living standards of the poor have risen. Conversely, a natural catastrophe which hurts absolutely everyone will decrease poverty if the rich are proportionately the most hurt. When used alone, relative poverty lines can thus be shown to drift the analysis toward the concept of relative inequality, and away from absoluteness of deprivation in any of the poverty measurement spaces defined above. Because of this, they are probably best used in conjunction with absolute living standard thresholds, at least when the aim is to capture both absolute deprivation in basic physiological capabilities and social exclusion and relative deprivation in more social capabilities.

The second alternative poverty line approach relies on the use of subjective information on the link between living standards and well-being. One source of information comes from interviews on what is perceived to be a sound poverty line, as gathered for instance in Goedhart, et al. (1977).

> *We would like to know which net family income would, in your circumstances, be the absolute minimum for you. That is to say, that you would not be able to make both ends meet if you earned less. (p.510)*

The information can also come from individuals' direct appraisal of their welfare, or by asking people directly whether they feel poor or not. The answers are subsequently regressed on the living standards of the respondents to predict at which living standard they could be expected to be exactly at the poverty line.

One difficulty with the subjective approach is the sensitivity of poverty line estimates to the formulation of interview questions. Perhaps a more fundamental disquieting output is the considerable variability in the answers provided, even within groups of relatively socio-economically homogeneous respondents. How ought we to interpret this variability? Is it that the link between living standards and well-being varies systematically within homogeneous groups of people? If so, then we should not attempt to use living standards or other direct or indirect indicators of well-being to classify the poor and the non-poor. Instead, we should perhaps take individuals at their word on whether or not they are poor. But this would alternatively raise important practical problems for the assessment and the implementation of public policy. Can public policy rely appropriately and confidently for its implementation on the provision of subjective information on the part of individuals?

V. Public Policy and Safety Nets

Poverty measurement plays a central role in the discussion of public policy and safety nets in particular. It is used, among other things, to identify the poor and the non-poor, to design optimal poverty relief schemes, to estimate the errors of exclusion and inclusion in the set of the poor (also known as Type I and Type II errors), and to assess the equity of poverty alleviation policy. How many of the poor, for instance, are excluded from safety net programs? Is it the poorest of the poor who benefit most from public policy? Would a different sort of poverty alleviation policy reduce deprivation further?

Poverty Measurement and Public Policy

An important example of the central role of poverty measurement in the setting of public policy is the optimal selection of safety net targeting indicators. The theory of optimal targeting suggests that it will commonly be best to target individuals on the basis of indicators that are as easily observable and as exogenous as possible, while being as correlated as possible with the true poverty status of the individuals. Indicators that are not readily observable by program administrators are of little practical value. Indicators that can be changed effortlessly by individuals will be distorted by the presence of the program and will lose their poverty-informative value. Whether available indicators are sufficiently correlated with the deprivation of individuals in a population is given by a poverty profile. The value of this profile will naturally be highly dependent on the particular assumptions and the approach used to measure poverty.

Estimation of the errors of inclusion and exclusion of the poor is also a product of poverty profiling and measurement. These errors are central in the trade-off involved in choosing between broad coverage of the population, at relatively low administrative and efficiency costs, and a narrower coverage, with more generous forms of support, for fewer beneficiaries. However, as Van de Walle (1998) puts it, narrower coverage of the population, with presumably smaller errors of inclusion of the non-poor, does not inevitably lead to more equitable treatment of the poor.

13

Concentrating solely on errors of leakage to the non-poor can lead to policies which have weak coverage of the poor. (Van de Walle 1998, p.366).

The terms of this trade-off are again given by a poverty assessment exercise.

Another lesson of optimal redistribution theory is that it is ordinarily better to transfer resources from groups with a high level of average well-being to those with a lower one. What matters even more, however, is the distribution of well-being within each of the groups. For instance, equalizing mean well-being across groups does not usually eliminate poverty since there generally exist within-group inequalities. Even within the richer group, for instance, there normally will be found some deprived individuals, whom a rich-to-poor cross-group redistributive process would clearly not take out of poverty. The within- and between-group distribution of well-being that is required for devising an optimal redistributive scheme can be revealed by a comprehensive poverty profile.

Welfarist and Non-Welfarist Policy Implications

The distinction between the welfarist and non-welfarist approaches to poverty measurement often matters (implicitly or explicitly) for the assessment and the design of public policy. As described above, a welfarist approach holds that individuals are the best judges of their own well-being. It would thus in principle avoid making appraisals of well-being that conflict with the poor's views of their own situation. A typical example of a welfarist public policy would be the provision of adequate income-generating opportunities, leaving individuals to decide and reveal whether these opportunities are utility-maximizing, keeping in mind the other non-income-generating opportunities that are open to them.

A non-welfarist policy analyst would argue, however, that raising income opportunities is not necessarily the best policy option. This is partly because individuals are not necessarily best left to their own resolutions, at least in an intertemporal setting, for their educational and environmental choices, for instance. In other words, the poor's short-run preoccupations may harm their long-term self-interest. For example, individuals may choose not to attend skill-enhancing programs because they appear overly time-consuming in the short-run, and because they are not convinced or aware of their long-term benefits.

Hence, if left to themselves, the poor will not necessarily spend their income increase on functionings that basic-needs analysts would normally consider a priority, such as good nutrition and health. Thus, "basic needs" cannot be fulfilled only by the generation of private income, but may require significant amounts of targeted and in-kind public expenditures on areas such as education, public health, and the environment. This would be so even if the poor did not presently believe that these areas were deserving of public expenditures. Furthermore, social cohesion concerns are not well addressed by the maximization of private utility, and raising income opportunities will not fundamentally solve problems caused by adverse intra-household distributions of well-being, for instance.

An objection to the basic needs approach is that it is clearly paternalistic since it supposes that it is in the absolute interests of all to meet a set of often arbitrarily specified needs. Indeed, as emphasized above, non-welfarist approaches in general may use criteria for identifying and helping the poor that may conflict with the poor's views, preferences, and utility-maximizing options. For poverty alleviation purposes, this could go as far as enforced enrolment in

14

community development programs. This would not only conflict with the preferences of the poor, but would also clearly undermine their freedom to choose. Freedom to choose may, however, be one of the basic capabilities which contribute fundamentally to well-being.

A further example of the possible tension between welfarist and non-welfarist approaches to public policy comes from optimal taxation theory, which is linked to optimal poverty alleviation theory. In the tradition of classical microeconomics, which values leisure in the production choices and labor market decisions of individuals, pure welfarists would incorporate the utility of leisure in the overall utility function of workers, poor and non-poor alike. Classical optimal taxation theory then shows that giving a positive weight to such things as leisure suggests a generally lower benefit reduction rate on the income of the poor than otherwise. Taking into account abstract things such as the utility of leisure is less typical of the basic needs and functioning approaches. Such approaches would, therefore, usually be less reluctant to target program benefits more sharply on the poor, and exact steeper benefit reduction rates as income or well-being increases.

Relative to the pure welfarist approach, non-welfarist approaches are also typically less reluctant to impose utility-decreasing (or "workfare") costs as side effects of participation in poverty alleviation schemes. These side effects are in fact often observed in practice. For instance, it is well-known that public programs frequently imposed participation costs on benefit claimants. These are typically non-monetary costs. Such costs can be both physical and psychological: providing manual labour, spending energy, spending time away from home, sacrificing leisure and home production, finding information about application and eligibility conditions, corresponding and dealing with the benefit agency, standing in lines, keeping appointments, complying with application conditions, revealing personal information, feeling "stigma" or a sense of guilt, etc. Although non-monetary, these costs have a clear impact on participants' net utility from participating in the programs. When they are negatively correlated with unobserved (or difficult to observe) entitlement indicators, they can provide self-selection mechanisms that enhance the efficiency of poverty alleviation programs, for welfarists and non-welfarists alike. One unfortunate effect of these costs is, however, that many truly entitled and truly deserving individuals may shy away from the programs because of the costs they impose. Although program participation could raise their income and consumption above a money-metric poverty line, some individuals will prefer not to participate, revealing that they find apparent poverty utility preferable to program participation.

The breadth of the definition of functionings is clearly also important for the assessment and the design of public policy. For instance, public spending on education is often promoted on the basis of its impact on productivity and growth. But education can also be seen as a means to attain the functioning of literacy and participation in the community. This provides additional strong support for public expenditures on education. Analogous arguments also apply, for instance, to public expenditures on health, transportation, and the environment.

VI. Empirical Measurement of Poverty and Vulnerability

Poverty assessment is customarily carried out using data on households and individuals. These data can be administrative (i.e., stored in government files and records), they can come from censuses of the entire population, or (most commonly) they can be generated by probabilistic surveys on the characteristics and living conditions of a population of households.

Survey Issues

There are several aspects of the surveying process that are important for poverty assessment. First, there is the coverage of the survey: does it contain representative information on the entire population of interest, or just on some socio-economic subgroups? Whether the representativeness of the data is appropriate depends on the focus of the poverty assessment. A survey containing observations drawn exclusively from the cities of a particular country will be perfectly fine if the aim is to design poverty alleviation schemes within these cities; its representativeness will, however, be insufficient if the objective is to assess the allocation of resources between the country's urban and rural areas.

Then there is the sample frame of the survey. Surveys are usually divided into strata and clusters. Stratification ensures that information is obtained from each of a given number of areas within a population of interest. Population strata are often geographic and can represent, for instance, the different regions or provinces of a country. Clustering facilitates the interviewing process by concentrating sample observations within particular population subgroups or geographic areas. Strata are thus often divided into a number of different levels of clusters, representing, say, cities, villages, neighborhoods, or households. A complete listing of the clusters in each stratum is then used to select randomly within each stratum a particular number of clusters. The selected clusters can then be subjected to further stratification or clustering, and the process continues until the final sampling units (usually households or individuals) have been selected and interviewed.

Fundamental in the use of survey data is the role of the randomness of the information that is generated. Because households and individuals are not all systematically interviewed (unlike for the case of censuses), the information generated by the use of survey data will depend on the precise households and individuals that have been selected in a particular sample. A poverty assessment of a given population will then vary randomly across the various samples that can be selected from this same population. For that reason, poverty assessments carried out using survey data will be subject to so-called "sample errors," that is, to sampling variability. When generating population poverty assessments from sample survey data, it is therefore important to recognize and assess the statistical imprecision of the sampling results obtained.

By ensuring that a minimum number of observations is obtained from each of a number of strata, stratification decreases the extent of sampling errors. A similar effect is obtained by increasing the total size of the sample: the greater the number of households surveyed, the greater on average is the precision of the estimates obtained. Conversely, by bundling observations around common geographic or socio-economic indicators, clustering tends to increase the effect of sampling errors on poverty assessment exercises. The sampling frame of a survey also impacts on its ability to provide accurate information on certain population subgroups. For instance, if the clusters within a stratum represent regions, and between-region variability is large, it would not be reasonable to use the information generated by the selected regions to depict poverty in the other, non-selected regions.

Survey data are also fraught with measurement and other "non-sampling" errors. For instance, even though they have been selected for appearance in a sample, some households will not be interviewed, either because they cannot be reached or because they refuse to be interviewed. Such "non-response" will raise difficulties for poverty assessments if it is correlated with observable and non-observable household characteristics. Even if

16

interviewed, households will sometimes consistently misreport their characteristics and living conditions, either because of ignorance or because of self-interest. This tends to make poverty assessments built from survey data diverge systematically from the true (and unobserved) population poverty assessment that would be carried out if there were no sampling and non-sampling errors. Clearly, such a shortcoming can bias the understanding of poverty and the consequent design of public policy.

The empirical analysis of vulnerability and poverty dynamics is particularly "data demanding." In general, it requires longitudinal (or panel) surveys, which follow each other in time and interview the same final observational units. Because they link the same units across time, they contain more information than the transversal (or cross-sectional) surveys, and they are particularly useful for measuring vulnerability and for understanding poverty dynamics—in addition to facilitating the assessment of the temporal effects of public policy on well-being. It must be stressed, however, that measurement error problems render the analysis of vulnerability and mobility very difficult, and results must be interpreted with caution.

Income versus Consumption

It is frequently argued that consumption is better suited than income as an indicator of living standard, at least in many developing countries. One reason is that consumption is believed to vary more smoothly than income, both within any given year and across the life cycle. Income is notoriously subject to seasonal variability, particularly in developing countries, whereas consumption tends to be less variable. Life-cycle theories predict that individuals will try to smooth their consumption across their low- and high-income years (in order to equalize their "marginal utility of consumption" across time), through appropriate borrowing and saving. In practice, however, consumption smoothing is far from perfect, in part due to imperfect access to commodity and credit markets and to difficulties in estimating precisely one's "permanent" or average life-cycle income.

For the non-welfarist interested in outcomes and functionings, consumption is also preferred over income because it is deemed to be a more "direct" indicator of achievements and fulfillment of basic needs. A caveat is, however, that consumption is also an outcome of individual free choice, an outcome which may differ across individuals with the same income and ability to consume, just like the actual functionings vary across people with the same capability sets. For a given capability to spend, some individuals may be left to consume less (or little), choosing instead to give to charity, to vow poverty, or to save in order to give important bequests to their children.

Consumption is also held to be more easily observable and measurable than income in developing countries (although this is not always the case). This is not to say that consumption is easy to measure correctly. For one thing, consumption does not equal expenditures. Unlike expenditures, consumption includes the value of self-produced goods. The value of these goods is not easily assessed, since they have not been transacted in a market. Distinguishing consumption expenditures from investment expenditures is very difficult, but failure to do so properly can lead to double-counting in the consumption measure. For instance, a $1 expenditure on education or machinery should not be counted as current consumption if the returns and the utility of such expenditure will only accrue later in the form of higher future earnings. Similarly, the value of the services provided by those durable goods owned by individuals ought also to enter a complete consumption indicator,

but the cost of these durable goods should not enter entirely the consumption aggregate at the time at which the good is purchased. An important example of this is owner-occupied housing. Again, estimating the service value of durable goods is not easily done. Further difficulties arise from the assessment of the value of various non-market goods and services, such as those provided freely by the government, and intangible benefits such as the quality of the environment, security, peace, and so on.

Price Variability

Whether it is income or consumption that is measured and compared, an important issue is how to account for the variability of prices across space and time. Conceptually, this also includes variability in quality and in quantity constraints. Failure to account for such variability can distort comparisons of well-being across time and space. In Ecuador, for instance (see Hentschel and Lanjouw 1996), and in many other countries, some households have free access to water, and tend to consume relatively large quantities of it with zero water expenditure. Others (often peri-urban dwellers) need to purchase water from private vendors and consequently consume a lower quantity of it at necessarily higher total expenditures. Ranking of households according to water expenditures could wrongly suggest that those who need to buy water are richer and derive greater utility from water consumption (since they spend more on it).

Microeconomic theory suggests that we may wish to account for price variability by comparing *real* as opposed to *nominal* consumption (or income). This can first be done by estimating the parameters of the indirect utility function of the economy's consumers. These parameters identify the ordinal preferences of the consumer. Inversion of the indirect utility function yields an equivalent (or real) consumption function, which indicates how much consumption at reference prices is needed to be equivalent (or to generate the same utility) as the consumption observed at current prices.

A second procedure deflates by a cost-of-living index the level of nominal consumption. This cost-of-living index is often taken to be the consumer price indices routinely computed by national statistical agencies. These consumer price indices ordinarily vary across regions and time, but not across levels of living standards (namely, across the poor and the non-poor). In some circumstances (i.e., for homothetic utility functions and when consumer preferences are identical), the above two procedures are equivalent. In general, however, they are not the same.

One important implication of this for poverty measurement is that the true cost-of-living index would normally be different across the poor and the rich. Using the same price index for the two groups may distort comparisons of well-being. An example is the effect of an increase in the price of food on economic well-being. Since the share of food in total consumption is habitually higher for the poor than for the rich, this increase should hurt disproportionately more the more deprived. Deflating nominal consumption by the same index for the entire population will, however, suggest that the impact of the food price increase is shared proportionately by all.

In few developing countries, however, are consumer price indices available or sufficiently disaggregated spatially. The alternative is then to produce different poverty lines for different regions (based on the same consumption basket) or construct food price indices.

18

In both cases, the analyst would usually be using regional price information derived from LSMS-style survey data. The resulting indices would then be interpreted as cost-of-living indices, and would help correct for spatial price variation.

Household Heterogeneity

A fundamental problem arises when comparing the needs of individuals who live in households of differing sizes and composition. Such comparisons commonly involve the use of equivalence scales. With these scales, a household of a particular size and composition is said to be comparable to a household of a particular number of "reference" or "equivalent" adults. Strategies for the estimation of equivalence scales are all contingent on the choice of comparable indicators of well-being. All such indicators are, however, intrinsically arbitrary. A popular example is food share in total consumption: at equal household food shares, individuals of various household types are deemed to be equally well-off. But, at equal well-being, one household type can well choose a food share that differs from that of the other household types. This would be the case, for instance, for households of smaller sizes for which it would make "sense" to spend more on food than on goods for which economies of scale are larger, such as housing.

Another difficulty arises when household size and composition are the result of a deliberate free choice. It may be argued, for instance, that a couple that elects freely to have a child cannot perceive this increase in household size to be utility-decreasing. This would be so even if the household's total consumption remained unchanged after the birth of the child (or even if it fell), despite the fact that most poverty analysts would judge this birth to increase household "needs." A final difficulty lies in the fact that the intra-household decision-making process can influence adversely the allocation of resources across household members, and thereby lead to wrong inferences of comparative needs. This is the case, for instance, when more is spent on boys than girls, not because of differential needs, but because of differential preferences on the part of the household decision-maker. Using these observed preferences to estimate household needs would then underestimate on average the level of deprivation experienced by girls and their families, since it would be wrongly assumed that girls are less "needy." A similar difficulty of interpretation arises when the household decision-maker is a man, and the consumption of his spouse is observed to be smaller than his own.

An additional problem in measuring individual living standards using survey data comes from the presence of intrahousehold inequality. The final unit of observation in surveys is customarily the household. Little information is typically provided on the intrahousehold allocation of well-being (e.g., of the individual benefits of total household consumption). Because of this, the usual procedure is to assume that the adult-equivalent consumption (once computed) is enjoyed identically by all household members. This, however, is at best an approximation of the true distribution of economic well-being in a household. If the nature of intrahousehold decision-making leads to important disparities in well-being across individuals, assuming equal sharing will significantly underestimate inequality and aggregate poverty. Not being able to account for intrahousehold inequities will also have important implications for profiling the poor, and also for the design of public policy. For instance, a poverty assessment that correctly showed the deprivation effects of unequal sharing within households could indicate that it would be relatively inefficient to target support at the level of the entire household, without taking into account how the targeted resources would

subsequently be allocated within the household. Instead, it might be better to design public policy such as to self-select the least privileged within the households, in the form of specific in-kind transfers or specially designed incentive schemes.

A final and related difficulty concerns who we are counting in aggregating poverty: is it individuals or households? Although this distinction is fundamental, it is often surprisingly hidden in applied poverty profile and poverty measurement papers. The distinction matters since there is habitually a strong positive correlation between household size and a household's poverty status. Expressed differently, household poverty is found disproportionately among the larger households. Because of this, counting households instead of individuals will typically underestimate significantly the true proportion of individuals in poverty.

VII. Comparing Poverty

Making poverty comparisons is essential to determine whether poverty has changed across time, or how it compares across countries, regions, or socio-economic groups. Poverty comparisons are also essential for designing public policy, and for assessing its effects on poverty. They may be used, for instance, to judge whether and by how much a public safety net reduces poverty and whether reforming its structure could further alleviate poverty.

Cardinal versus Ordinal Comparisons

There are two types of poverty comparisons, cardinal and ordinal. Cardinal poverty comparisons simply involve differences in numerical poverty estimates. Numerical poverty estimates attach a single number to the extent of poverty in a population, for instance, 40 percent or $200 per capita. These estimates are valuable when a precise number must be attached to the extent of poverty in a distribution of well-being. Cardinal poverty estimates require specific and precise assumptions, such as the nature of the poverty index, the definition of the indicator of well-being, the value of the poverty line, and how that poverty line varies exactly across household types, regions and time. Once this information is provided, cardinal poverty estimates can tell, for instance, that 30 percent of the individuals in a population used to have consumption below the poverty line, but that a recently-introduced public safety net has decreased that proportion to 25 percent. Cardinal poverty estimates can also be used to carry out a money-metric cost-benefit analysis of the effects of safety nets. Thus, if the above safety net involved yearly expenditures of $500 million, then we would know immediately that a 1 percent fall in the proportion of the poor would seem to cost the government on average $100 million. That amount could then be compared to the poverty alleviation cost of other forms of government policy.

The main advantage of cardinal poverty estimates is their ease of communication, their ease of manipulation, and their (apparent) lack of ambiguity. Government officials and the media often want the results of poverty comparisons to be produced in straightforward and precise terms, and can feel annoyed when this is not possible. Cardinal poverty estimates are, however, necessarily (and often highly) sensitive to the choice of a number of arbitrary measurement assumptions. It is clear, for example, that choosing a different poverty line will almost always change the estimated numerical value of any index of poverty. The elasticity of the poverty headcount index to the poverty line is, for example, almost always significantly larger than 1. This implies that a variation of 10 percent in the poverty line will change by more

than 10 percent the estimated proportion of the poor in the population; this is a substantial impact for those interested in poverty alleviation, especially since poverty lines are rarely convincingly bounded within a narrow confidence interval. Another source of numerical variability comes from the choice of the form of the poverty index. Many procedures have been proposed to aggregate numerically the poverty of individuals. Depending on the chosen procedure, numerical estimates of poverty will appear large or low. As we will see later, for instance, the estimation of a "socially representative poverty gap" will rest particularly on the weight given to the more deprived among the poor. There is little objective guidance in choosing that weight; the greater that weight, however, the greater the estimated socially representative poverty gap, and the greater the numerical estimate of poverty.

To see this better, consider the hypothetical example of table 1. The top four lines in the table show the standards of livings in two distributions, A and B. Thus, distribution A contains three standards of living of 4, 11, and 20 respectively. The bottom four lines of the table show the value of two different popular indices of poverty, the headcount and the average poverty gap indices, at two alternative poverty lines, 5 and 10. As we will see more formally in Section 0, the poverty headcount gives the proportion of individuals in a population whose standard of living falls below a poverty line. At a poverty line of 5, there is only one person in poverty in distribution A, and the headcount is thus equal to $1/3 = 0.33$. The average poverty gap index is the sum of the distance of the poor's standards of living from the poverty line, divided by the number of people in the population. For instance, at a poverty line of 10, there are two people in poverty in B, and the sum of their distance from the poverty line is $(10 - 6) + (10 - 9) = 5$. Divided by 3, this gives 1.66 as the average poverty gap in B for a poverty line of 10. The last column of table 1 gives the poverty ranking of the two distributions according to the different choices of poverty lines and poverty indices. At a poverty line of 5, distribution A has clearly more poverty than B, but the ranking is spectacularly reversed if we consider instead the headcount and a poverty line of 10. The ranking changes again if we use the same poverty line of 10 but now focus on the average poverty gap as the poverty index. Clearly, here, ranking A and B can be quite sensitive to the precise choice of measurement assumptions.

Table 1. Difficulties in Comparing Poverty

		Incomes under Distribution A	Incomes under Distribution B	
		4	6	
		11	9	
		20	20	
Poverty line	Poverty index	Index for Distribution A	Index for Distribution B	Poverty ranking
5	Headcount	0.33	0	$P_A > P_B$
5	Average poverty gap	0.33	0	$P_A > P_B$
10	Headcount	0.33	0.66	$P_A < P_B$
10	Average poverty gap	2	1.66	$P_A > P_B$

Ordinal comparisons, on the other hand, do not attempt to put a precise numerical value on the extent of poverty. They only try to *rank* poverty across two distributions, indicating whether it is higher or lower in the first than in the second. Ordinal comparisons of poverty do not, therefore, provide precise numerical data to compare with metric indicators of other

aspects or effects of government policy, such as its administrative or efficiency cost. This is their main defect. They can, however, be highly robust to the choice of measurement assumptions, since they will sometimes be valid for wide ranges of such assumptions. When the problem is simply of resolving which of two policies will better alleviate poverty, or determining which of two distributions has the most poverty, ordinal comparisons can be sufficiently informative, that is, cardinal estimates will not be needed. In that case, ordinal comparisons will also be sufficiently convincing. For instance, we will see later in Section 0 that we can order robustly distributions A and B in table 1 for all "distribution-sensitive" poverty indices and for any choice of poverty line.

A focus on ordinal comparisons has two major advantages. First, it saves most of the considerable energy and time often spent on choosing poverty lines and poverty indices. This includes avoiding the difficult debate on the choice of appropriate theoretical and econometric methods for estimating poverty lines. It also enables the poverty analyst to escape arguing on the relative merits and properties of the many poverty indices that have been proposed in the scientific literature. This is because ordinal poverty comparisons do not require that the precision of numerical poverty estimates be validated; it is simply their ordinal ranking across policies or distributions of well-being that is important, and for this, it is not required that the poverty estimates be precisely known.

VIII. Aggregate Poverty Indices

Two approaches have been used to devise cardinal aggregate indices of poverty using unidimensional indicators of well-being, typically consumption expenditures or income. The first uses the concept of the equally distributed equivalent (EDE) living standard of a society where living standards have been censored at the poverty line, and compares it to the poverty line. The second combines living standards and the poverty line into poverty gaps, and aggregates them in social-welfare-like functions to assess overall poverty. We look at these two approaches in turn.

EDE Poverty Indices

For the EDE approach to building poverty indices, we simply use the distribution of living standards, which we denote as y_i. Say that there are N individuals in a population. Since, for poverty comparisons, we want to focus on the living standards that fall below the poverty line (the so-called "focus axiom"), the living standards are censored (or "cut off") at the poverty line z, to give y_i^*. Mathematically, we have $y_i^* = \min(y_i, z)$. The censored living standards are then aggregated using one of the many social welfare functions that have been proposed in the literature. The EDE living standard is the equally distributed living standard which yields the same value for the social welfare function as the distribution of the y_i^*. A poverty index is obtained by taking the difference between the poverty line and the EDE censored living standard. This poverty index can then be interpreted as the "socially representative" or EDE poverty gap.

Examples of such EDE poverty indices include (a transformation of) the Clark, Hemming and Ulph (CHU) second class of poverty indices, which are closely related to the well-known Atkinson social welfare functions and inequality indices. When the inequality aversion parameter of the Atkinson functions is set to 1, the CHU poverty index is also the EDE poverty gap corresponding to the early Watts (1968) poverty index. Values of the

inequality aversion parameter between 0 and 1 for the CHU indices also correspond to the EDE poverty gap of the class of poverty indices proposed by Chakravarty (1983).

The class of S-Gini indices of poverty is obtained by using the S-Gini social welfare function (see Donaldson and Weymark 1980 and Yitzhaki 1983). For the special case of the standard Gini social welfare function (which is linked to the traditional and well-known Gini index of inequality), the procedure yields the Thon-Chakravarty-Shorrocks index of poverty. This index can also be more simply referred to as the "Gini" index of poverty, of which the seminal Sen (1976) poverty index is a particular transformation.

All of the above indices are decreasing in living standards, that is, they fall when the living standard of a poor person increases. They are also said to be distribution sensitive (except for Sen's index): a transfer from a richer person (though possibly poor) to a poorer one decreases aggregate poverty.

Poverty Gap Indices

The most popular approach to aggregating poverty has, however, focused attention on the distribution of poverty gaps, $z - y_i^*$. After the distribution of poverty gaps has been computed, use is made of aggregating functions analogous to those mentioned above. The most popular class of poverty gap indices is known as the Foster-Greer-Thorbecke (1984) (FGT) class, and is defined as:

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

where $\alpha > 0$. When $\alpha = 0$, the FGT index gives the simplest and most common example of a poverty index. $P_0(z)$ is called the headcount ratio, and is simply the proportion of the poor (those with a positive poverty gap) in a population. The next simplest and most commonly used index is given by $\alpha = 1$. $P_1(z)$ is the sum of the poverty gaps divided by N. It is therefore the average poverty gap, and it equals the average shortfall of living standards from the poverty line in the population. In the event that a government could perfectly target the poor, and that the poor's living standard y_i were not affected by this targeting, $P_1(z)$ would indicate the level of per capita expenditures that would be necessary to eliminate poverty. For $\alpha = 2$, we obtain the popular sum of the squared poverty gaps, an index that is "distribution sensitive," just like all of the FGT indices for which $\alpha \geq 1$.

Other poverty gap indices can be easily proposed, simply by using other aggregating functions of poverty gaps that obey some of the desirable axioms (such as that of being increasing and distribution sensitive in poverty gaps) discussed in the literature.

Group-Decomposable Poverty Indices

Much of the literature on the construction of poverty indices has focused on whether indices are decomposable across population subgroups. This has led to the identification of a subgroup of poverty indices known as the "class of decomposable poverty indices." These indices have the property of being expressible as a weighted sum (more generally, as a separable function) of the same poverty indices assessed within population subgroups. They most commonly include the FGT and the Chakravarty classes of indices, as well as the Watts index.

Let the population be divided into K mutually exclusive population subgroups, where φ_k is the share of the population found in subgroup k. For the FGT index, we then have that:

$$P_\alpha(z) = \sum_{k=1}^{K} \varphi_k P_k^\alpha(z)$$

where $P_k^\alpha(z)$ is the FGT poverty index of subgroup k. The Watts and Chakravarty indices are expressible as a sum of the poverty indices of each subgroup in exactly the same way as for the FGT indices.

> *Subgroup decomposability implies that an improvement in one of the subgroups will necessarily improve aggregate poverty if the living standards in the other groups have not changed. It will also mean that the design of social safety nets and benefit targeting within any given group can be done independently of the distribution of living standards in the other groups. This enables targeting to be done in a decentralized manner: only the characteristics of a relevant population matter for the exercise. If targeting · succeeds in decreasing poverty at a local level, then it must succeed also at the aggregate level.*

Subgroup decomposability is therefore useful, although it is certainly *not imperative* for poverty analysis. In particular, it must be admitted that it is not because an index property facilitates poverty profiling and the analysis of the comparative advantages of various forms of targeting that this property is ethically fine. Among other things, imposing the decomposability and additivity property can mean losing some important ethical aspects to the aggregation of poverty. In that context, Ravallion (1994) notes that when measuring poverty "one possible objection to additivity is that it attaches no weight to one aspect of a poverty profile: the inequality between sub-groups in the extent of poverty." This can be an important flaw if considerations of between-group relative deprivation are significant.

IX. Poverty Quantile Curves

It is generally informative to portray the whole distribution of poverty gaps on a simple graph, in a way that shows both the incidence and the inequality of the deprivation in living standards. Particularly useful are the *poverty gap curves*. To see how they can be computed, rank observations from the poorest to the richest in terms of income, such that they take increasingly high values as ranks increase: $y_1 \leq y_2 < \ldots \leq y_N$. Let p_i represent the rank of observation i in the population, lying between 0 and 1. Since there are N individuals, the rank of observation i is $p_i = i/N$, and y_i is called the $100 \cdot p_i \%$ income quantile. Recall that poverty gaps are denoted by $z - y^*$. The $100 \cdot p_i \%$ poverty gap quantile is then given by $g(p_i, z) = z - y_i^*$. When the quantiles of poverty gaps are graphed against values of p_i, we obtain the poverty gap curve (see figure 1). The poverty gap curve shows the "intensity of poverty" felt at each rank in the population. The curve naturally decreases with the rank p in the population, and reaches zero at the value of p equal to the headcount ratio. The integral under the curve gives the average poverty gap, and its steepness, the degree of inequality in the distribution of poverty gaps.

Figure 1: Poverty Gap Curve

Source: Author.

Another quantile-based curve that is graphically informative and that is useful for the measurement and comparison of poverty is called the Cumulative Poverty Gap (CPG) curve.[2] It is defined as:

$$G(p_i, z) = \frac{1}{N} \sum_{j=1}^{i} (z - y_i^*)$$

The CPG curve cumulates the poverty gaps of the bottom p_i proportion of the population, as shown in figure 2. The CPG curve enjoys for poverty comparisons the same interest as the Lorenz and Generalized Lorenz curves for inequality and social welfare analysis. Its slope at p_i shows the poverty gap quantile, $g(p_i, z) = z - y_i^*$. Its distance from the line of perfect equality of poverty gaps displays the inequality of poverty gaps among the population. $G(p_i = 1; z)$ equals the average poverty gap, and the horizontal value at which $G(p_i; z)$ becomes horizontal yields the poverty headcount. When weighted by 2, the area underneath the CPG curve generates the Gini poverty index.

---

[2] It is also sometimes referred to as the inverse generalized Lorenz curve, the "TIP" curve, or the poverty profile curve.

Figure 2: CPG Curve

Figure 2: CPG Curve

G(p,z)

CPG curve   Line of perfect equality
              of poverty gaps

P₁(z) → $P_1(z)$

$P_0(z)$ 1   p
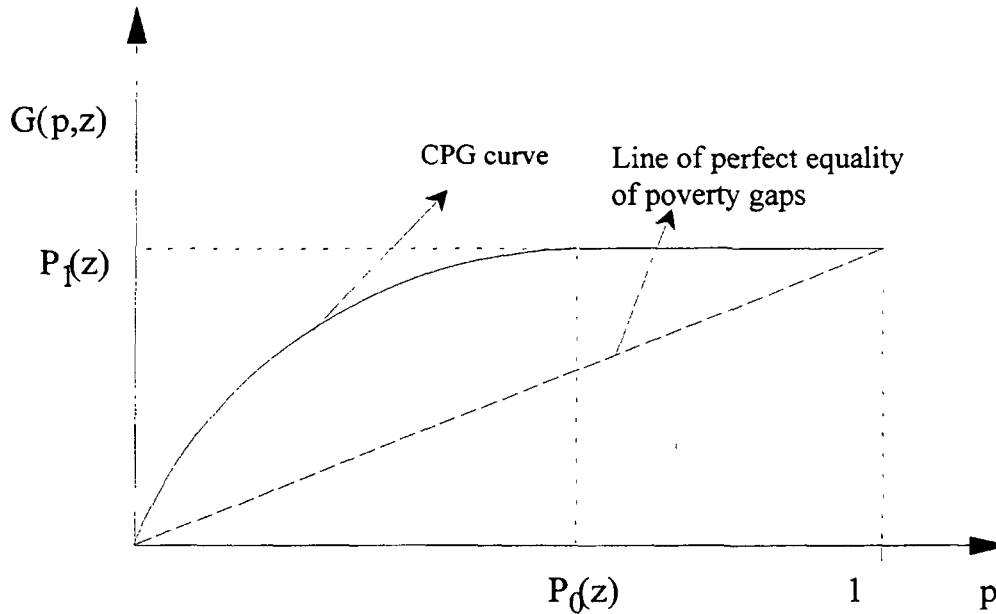
Source: Author

X. Poverty Dominance

The main reason for carrying out analyses of poverty dominance is that comparisons of poverty across time, regions, socio-demographic groups, or fiscal regimes (for instance) may be sensitive to the choice of the poverty line and to the choice of the poverty index. This is problematic since a different choice of poverty index or poverty line could reverse an earlier conclusion that poverty is greater in region A than in region B, or that poverty will decrease following the introduction of a particular fiscal policy or macroeconomic adjustment program. Such sensitivity must be checked for one to have some confidence that a poverty ordering is robust to the choice of a poverty line or of a poverty index. Another reason is that unknown errors in measuring well-being will necessarily affect cardinal poverty estimates; under some assumptions (admittedly restrictive), such errors will not contaminate ordinal poverty comparisons.

In essence, testing poverty dominance allows one to secure poverty comparisons that necessarily hold for groups (or classes) of poverty indices, as well as for ranges of poverty lines. These classes are defined for specific orders s of stochastic dominance. The first-order class of poverty indices regroups all poverty indices that weakly decrease when the living standard of someone in the population increases. By "weakly decrease," we mean that the poverty index will never increase following a rise in someone's living standard, and will sometimes decrease if the person involved initially had a living standard below the poverty line. These poverty indices have properties that are analogous to those of Paretian social welfare functions: other things being equal, the larger the levels of the individual living standards, the better off is society (the lower is poverty). The second-order class of poverty indices contains those indices (among the first-order class of indices) that have a greater ethical preference for the poorer among the poor. Mathematically, these indices are convex in

26

living standards: all other things being equal, the more equal the distribution of living standards among the poor, the lower the level of poverty. The indices thus display a preference for equality of living standards. If a transfer from a poor to a poorer person takes place without reversing the ranks of the two individuals, the indices in the class of second-order indices will never increase, and will sometimes fall. This equality-preferring property is analogous to the Pigou-Dalton principle of transfer for social welfare functions (social welfare *increases* when an equalizing transfer of living standards takes place). They are therefore "distribution-sensitive." All of the indices that belong to the second-order class of indices also belong to the first-order class.

To understand the third-order class of poverty indices, imagine four levels of living standard, for individuals 1, 2, 3, and 4, such that $y_2 - y_1 = y_4 - y_3 > 0$ and $y_1 < y_3$. Let a marginal transfer of \$1 of living standards be made from individual 2 to individual 1 (an equalizing transfer) at the same time as an identical \$1 is transferred from individual 3 to individual 4 (a disequalizing transfer). This is called in the literature a "favorable composite transfer." Note that the equalizing transfer is made lower down in the distribution of living standards than the disequalizing transfer. This can be seen by the fact that the recipient of the first transfer, 1, has a lower standard of living than the donor of the second transfer, 3, since $y_3 > y_1$.

> *There are often sound ethical reasons to be socially more sensitive to what happens toward the bottom of the distribution of living standards than higher up in it. We may thus be less concerned about the "bad" disequalizing transfer higher up in the distribution of living standards than we are pleased about the "good" equalizing transfer lower down. Second-order poverty indices which exhibit this property by decreasing when a favorable composite transfer is effected are said to belong also to the third-order class of poverty indices, and to obey the "transfer-sensitivity" principle. Mathematically, for these poverty indices we require their second-order derivative to be decreasing with income.*

We can, if we wish, define subsequent classes of poverty indices in an analogous manner. As the order *s* of the class of poverty indices increases, the indices become more and more sensitive to the distribution of living standards among the poorest. At the limit, as *s* becomes very large, only the living standard of the poorest individual matters in comparing poverty across two distributions.

A number of well-known poverty indices fit into some of the classes defined above. The headcount index belongs only to the first class. The average poverty gap belongs to the first and to the second, as do the Watts index, the Chakravarty and the CHU indices, and the Gini index of poverty. The squared-poverty-gap index belongs to the first three.

To check whether poverty in *A* is greater than in *B* for all indices that are members of any one of these classes, there exist two approaches: a primal approach and a dual approach. We look at them in turn.

Primal Approach to Poverty Dominance

We are interested in whether we may assert confidently that poverty in a distribution A is larger than poverty in a distribution B for all of the poverty indices belonging to one of the classes of poverty indices defined above and for a range of possible poverty lines. Checking for this in the primal approach makes use of dominance curves $D^s(z)$ for orders of dominance $s = 1, 2, 3, \ldots$ There exists a useful link between the dominance curves and the popular FGT indices, a link that greatly facilitates the computation of $D^s(z)$. Indeed, we can show that

$$D^s(z) = c.P_{s-1}(z)$$

where $c = 1/(s-1)!$ is a constant that can be safely ignored in the use of dominance curves. Therefore, to compute the dominance curve of order s, we only need to compute the FGT index at $\alpha = s - 1$. The first-order dominance curve is then given by the headcount index, the second-order curve by the average poverty gap index, the third-order curve by the squared-poverty-gap index, and so on.

Say that the range of admissible poverty lines is bounded upwards by z_{max}. To check for poverty dominance, whatever the order of dominance s, the approach is simple and always the same. For first-order poverty dominance, we need to check whether:

$$D_A^I(z) \geq D_B^I(z) \text{ for all } z \in [0, z_{max}].$$

For first-order dominance, we therefore require the poverty headcount in A to be always (weakly) larger than the headcount in B, for all of the poverty lines between 0 and z_{max}. It is a relatively stringent condition. If, however, it is found to hold in practice, then a very robust poverty ordering is obtained: we can then unambiguously say that poverty is higher in A than in B for *all* of the poverty lines between 0 and z_{max} and for *all* of the first-order class of poverty indices. These indices include all those which are weakly decreasing in living standards. Since almost all of the poverty indices that have been proposed obey this restriction, this is a very powerful conclusion indeed.

For second-order poverty dominance, we require that the average poverty gap in A be always larger than the average poverty gap in B, for all of the poverty lines between 0 and z_{max}:

$$D_A^2(z) \geq D_B^2 \text{ for all } z \in [0, z_{max}].$$

This is a less stringent condition than the first-order poverty dominance condition, since when first-order dominance over $[0, z_{max}]$ holds, then second-order dominance over $[0, z]$ must also hold, but not necessarily the converse. If second-order poverty dominance is found to hold in practice, then a rather robust poverty ordering is obtained: we can unambiguously say that poverty is higher in A than in B for *all* of the poverty lines between $[0, z_{max}]$ and for all of the second-order class of poverty indices. Most of the indices found in the literature fall into that category, a major exception being the headcount and the Sen indices. In fact, a comparison of distributions A and B in table 1 shows that this condition is obeyed for any choice of z_{max}. Hence, saying that A has more poverty than B in that table is quite a robust statement, since it is valid for all distribution-sensitive poverty indices (the headcount is not distribution-sensitive, hence it does not always indicate more poverty in A) and for any choice of poverty line. As mentioned, second-order poverty dominance is a less stringent

criterion than first-order dominance to check in practice. The price of this, however, is that the set of indices over which poverty dominance is checked is smaller for second-order dominance than for first-order dominance.

We can repeat this process for any arbitrary order of dominance. We will find more poverty in A than in B for *all* of the poverty lines between 0 and z_{max} and for all of the s-order poverty indices if

$$D_A^S(z) \geq D_B^S \text{ for all } z \in [0, z_{max}].$$

This is illustrated in figure 3, where dominance holds until z_{max}, but would not hold if z_{max} exceeded z^*. Checking poverty dominance is thus conceptually straightforward. For first-order dominance, we use what has been termed "the poverty incidence curve," which is the headcount index as a function of the range of poverty lines $[0, z_{max}]$. For second-order dominance, we use the "poverty deficit curve," which is the average poverty gap, again as a function of the range of poverty lines $[0, z_{max}]$. Third-order dominance makes use of the poverty severity curve, or the squared-poverty-gap index for poverty lines between 0 and z_{max}.

Figure 3: Poverty Dominance Curves

Source: Author.

Dual Approach to Poverty Dominance

There also exists a dual approach to testing first-order and second-order poverty dominance, which is sometimes called a percentile or quantile approach. Whereas the primal approach makes use of curves that focus on the population's living standards below varying poverty lines (0 to z_{max}), the dual approach makes use of curves that focus on the population at (or

below) a particular rank in the population. As discussed above, these curves have interesting graphical properties, which make their use instructive in checking poverty dominance.

For first-order dominance, we need to check that $g_A(p, z_{max}) \geq g_B(p, z_{max})$ for all ranks p. This requires poverty gaps to be nowhere lower in A than in B, whatever the ranks p considered. We can show that this is equivalent to the primal first-order poverty dominance condition described above. Therefore, when $g_A(p, z_{max}) \geq g_B(p, z_{max})$ for all p, poverty is higher in A than in B for *all* of the poverty lines between 0 and z_{max} and for *all* of the first-order class of poverty indices.

The dual second-order dominance condition is that $G_A(p, z_{max}) \geq G_B(p, z_{max})$ for all ranks p. This requires the CPG curve to be higher in A than in B. Again, we can show that this condition is equivalent to the primal second-order poverty dominance condition. When it is respected, poverty is therefore higher in A than in B for *all* of the poverty lines between 0 and z_{max} and for *all* of the second-order class of poverty indices.

References

The word *processed* describes informally reproduced works that may not be commonly available through libraries.

Alwang, Jeffrey, and Paul B. Siegel. 2000. "Towards Operational Definitions and Measures of Vulnerability: A Review of the Literature from Different Disciplines." World Bank, Human Development Network, Social Protection Unit. Washington, D.C.

Chakravarty, S.R. 1983. "A New Index of Poverty." *Mathematical Social Sciences* 6(July): 307–13.

Donaldson, D., and J.A. Weymark. 1980. "A Single Parameter Generalization of the Gini Indices of Inequality." *Journal of Economic Theory* vol. 22: 67–86.

Foster, J.E., J. Greer, and E. Thorbecke. 1984. "A Class of Decomposable Poverty Measures." *Econometrica* 52(3): 761–76.

Glewwe, P., and G. Hall. 1998. "Are Some Groups More Vulnerable to Macroeconomic Shocks than Others? Hypothesis Tests Based on Panel Data from Peru." *Journal of Development Economics* vol. 56: 181–206.

Goedhart, T., V. Halberstadt, A. Kapteyn, and B. Van Praag. 1977. "The Poverty Line: Concept and Measurement." *Journal of Human Resources* vol. 4: 503–20.

Hentschel, Jesko, and Peter Lanjouw. 1996. "Constructing an Indicator of Consumption for the Analysis of Poverty." Living Standards Measurement Study Working Paper No. 124. World Bank, Washington, D.C.

Jalan, Jyotsna, and Martin Ravallion. 1997. "Consumption Variability and Rural Poverty in Post-Reform China." World Bank, Processed.

Orshansky, M. 1965. "Counting the Poor: Another Look at the Poverty Profile." *Social Security Bulletin*(June): 3–29.

Ravallion, M. 1994. *Poverty Comparisons.* Switzerland: Harwood Academic Publishers. Rodgers, G., C. Gore, and J.B. Figueiredo, eds. 1995. *Social Exclusion: Rhetoric, Reality, Responses.* Geneva: International Labour Organization (International Institute for Labour Studies).

Rowntree, S. 1901. *Poverty, a Study of Town Life.* London: MacMillan.

Saunders, Peter. 1994. *Welfare and Inequality: National and International Perspectives on the Australian Welfare State.* Cambridge: Cambridge University Press.

Sen, A.K. 1976., "Poverty: An Ordinal Approach to Measurement." *Econometrica* 44(2): 219–31.

_____. 1981. *Poverty and Famine: An Essay on Entitlement and Deprivation.* Oxford: Oxford University Press.

_____. 1983. "Poor, Relatively Speaking." *Oxford Economic Papers* 35(March): 153–69.

_____. 1997. *Inequality Reexamined.* Cambridge, Mass.: Harvard University Press.

Silver, H. 1994. "Social Exclusion and Social Solidarity: Three Paradigms." *International Labour Review* 133(5-6): 531–76.

Smith, Adam. 1776. *An inquiry into the nature and causes of the wealth of nations*, Home University Library, London.

Streeten, Paul, et al. 1981. *First Things First: Meeting Basic Human Needs in the Developing World.* Oxford: Oxford University Press.

Townsend, Peter. 1979. *Poverty in the United Kingdom: A Survey of Household Resources and Standards of Living.* Berkeley: University of California Press.

Van de Walle, Dominique. 1998. "Assessing the Welfare Impacts of Public Spending." *World Development* 26(3): 365–79.

Watts, H.W. 1968. "An Economic Definition of Poverty." In D.P. Moynihan, ed., *On Understanding Poverty.* New York: Basic Books.

World Bank (2001), *World Development Report 2000/2001, Attacking Poverty*, Oxford University Press

Yitzhaki, S. 1983. "On an Extension of the Gini Index." *International Economic Review.* 24: 617–28.

Social Protection Discussion Paper Series Titles

No.	Title

0233   Public Attitudes Matter: A Conceptual Frame for Accounting for Political Economy in Safety Nets and Social Assistance Policies
by Carol Graham

0232   Ensuring Access to Essential Services: Demand-Side Housing Subsidies
by Harold M. Katsura and Clare T. Romanik

0231   Gender and Risk in the Design of Social Protection Interventions
by Kene Ezemenari, Nazmul Chaudhury, and Janet Owens

0230   Vulnerability and Poverty Measurement Issues for Public Policy
by Jean-Yves Duclos

0229   Emerging from Ethnic Conflict: Challenges for Social Protection Design in Transition Countries
by Christian Bodewig

0228   Child Work in Zambia: A Comparative Study of Survey Instruments
by Niels-Hugo Blunch, Amit Dar, Lorenzo Guarcello, Scott Lyon, Amy Ritualo and Furio Rosati

0227   Strengthening Public Safety Nets from the Bottom Up
by Jonathan Morduch and Manohar Sharma

0226   Incentives and the Role of Institutions in Provision of Social Safety Nets
by Chris de Neubourg

0225   Food-Based Safety Nets and Related Programs
by Beatrice Lorge Rogers, Ph.D. and Jennifer Coates, M.S.

0224   Subsidies as a Social Safety Net: Effectiveness and Challenges
by Harold Alderman

0223   Assisting the Poor with Cash: Design and Implementation of Social Transfer Programs
by Steven R. Tabor

0222   Effectiveness of Lending for Vocational Education and Training: Lessons from World Bank Experience
by S. Canagarajah, A. Dar, R. Nording and D. Raju

0221   Participation of Children in Schooling and Labor Activities: A Review of Empirical Studies
by Amit Dar, Niels-Hugo Blunch and Bona Kim

<u>No.</u>	<u>Title</u>
0113	Regulating Private Pension Funds' Structure, Performance and Investments: Cross-country Evidence by P.S. Srinivas, Edward Whitehouse and Juan Yermo
0112	The World Bank and the Provision of Assistance to Redundant Workers: Experience with Enterprise Restructuring and Future Directions by Yi Chen
0111	Labor Markets in Transition Economies: Recent Developments and Future Challenges by Mansoora Rashid and Jan Rutkowski
0110	Operating Instruction Included: A Review of Social Investment Fund Operations Manuals by Juliana Weissman
0109	Risk and Vulnerability: The Forward Looking Role of Social Protection in a Globalizing World by Robert Holzmann
0108	Australia's Mandatory Retirement Saving Policy: A View from the New Millennium by Hazel Bateman and John Piggott
0107	Annuity Markets and Benefit Design in Multipillar Pension Schemes: Experience and Lessons from Four Latin American Countries by Robert Palacios and Rafael Rofman
0106	Guide for Task Teams on Procurement Procedures Used in Social Funds by Jorge A. Cavero Uriona
0105	Programmes Actifs Pour Le Marche Du Travail: Un Aperçu General Des Evidences Resultant Des Evaluations by Zafiris Tzannatos and Amit Dar
0104	Kazakhstan: An Ambitious Pension Reform by Emily S. Andrews
0103	Long-term Consequences of an Innovative Redundancy-retraining Project: The Austrian Steel Foundation by Rudolf Winter-Ebmer
0102	Community Based Targeting Mechanisms for Social Safety Nets by Jonathan Conning and Michael Kevane (available in Spanish)

Social Protection Discussion Paper Series Titles continued

No.	Title
0101	Disability and Work in Poland by Tom Hoopengardner
0024	Do Market Wages Influence Child Labor and Child Schooling? by Jackline Wahba
0023	Including the Most Vulnerable: Social Funds and People with Disabilities by Pamela Dudzik and Dinah McLeod
0022	Promoting Good Local Governance through Social Funds and Decentralization by Andrew Parker and Rodrigo Serrano
0021	Creating Partnerships with Working Children and Youth by Per Miljeteig
0020	Contractual Savings or Stock Market Development. Which Leads? by Mario Catalan, Gregorio Impavido and Alberto R. Musalem
0019	Pension Reform and Public Information in Poland by Agnieszka Chlon
0018	Worker Reallocation During Estonia's Transition to Market: How Efficient and How Equitable? by Milan Vodopivec
0017	How Poor are the Old? A Survey of Evidence from 44 Countries by Edward Whitehouse
0016	Administrative Charges for Funded Pensions: An International Comparison and Assessment by Edward Whitehouse
0015	The Pension System in Argentina: Six years after the Reform by Rafael Rofman
0014	Pension Systems in East Asia and the Pacific: Challenges and Opportunities by Robert Holzmann, Ian W. Mac Arthur and Yvonne Sin
0013	Survey of Disability Projects. The Experience of SHIA, Swedish International Aid for Solidarity and Humanity by Kaj Nordquist

Social Protection Discussion Paper Series Titles continued

Social Protection Discussion Paper Series Titles continued

No.   **Title**

* The papers below (No. 9801-9818 and 9901-9934) are no longer being printed, but are available for download from our website at www.worldbank.org/sp

9934   Helping the Poor Manage Risk Better: The Role of Social Funds
by Steen Jørgensen and Julie Van Domelen

9933   Coordinating Poverty Alleviation Programs with Regional and Local Governments: The Experience of the Chilean Social Fund – FOSIS
by Jorge C. Barrientos

9932   Poverty and Disability: A Survey of the Literature
by Ann Elwan

9931   Uncertainty About Children's Survival and Fertility: A Test Using Indian Microdata
by Vincenzo Atella and Furio Camillo Rosati

9930   Beneficiary Assessment Manual for Social Funds
by Lawrence F. Salmen

9929   Improving the Regulation and Supervision of Pension Funds: Are There Lessons from the Banking Sector?
by Roberto Rocha, Richard Hinz, and Joaquin Gutierrez

9928   Notional Accounts as a Pension Reform Strategy: An Evaluation
By Richard Disney

9927   Reform Options for Pay-As-You-Go Public Pension Systems
by Sheetal K. Chand and Albert Jaeger

9926   An Asset-Based Approach to Social Risk Management: A Conceptual Framework
by Paul Siegel and Jeffrey Alwang

9925   Migration from the Russian North During the Transition Period
by Timothy Heleniak

9924   Pension Plans and Retirement Incentives
by Richard Disney and Edward Whitehouse

9923   Shaping Pension Reform in Poland: Security Through Diversity
by Agnieszka Chlon, Marek Góra and Michal Rutkowski

Social Protection Discussion Paper Series Titles continued

No.	Title
9909	Russia's Social Protection Malaise: Key Reform Priorities as a Response to the Present Crisis edited by Michal Rutkowski
9908	Causalities Between Social Capital and Social Funds by Jesper Kammersgaard
9907	Collecting and Transferring Pension Contributions by Rafael Rofman and Gustavo Demarco
9906	Optimal Unemployment Insurance: A Guide to the Literature by Edi Karni
9905	The Effects of Legislative Change on Female Labour Supply: Marriage and Divorce, Child and Spousal Support, Property Division and Pension Splitting by Antony Dnes
9904	Social Protection as Social Risk Management: Conceptual Underpinnings for the Social Protection Sector Strategy Paper by Robert Holzmann and Steen Jorgensen (available in Russian)
9903	A Bundle of Joy or an Expensive Luxury: A Comparative Analysis of the Economic Environment for Family Formation in Western Europe by Pierella Paci
9902	World Bank Lending for Labor Markets: 1991 to 1998 by Amit Dar and Zafiris Tzannatos
9901	Active Labor Market Programs: A Review of the Evidence from Evaluations by Amit Dar and Zafiris Tzannatos
9818	Child Labor and School Enrollment in Thailand in the 1990s By Zafiris Tzannatos
9817	Supervising Mandatory Funded Pension Systems: Issues and Challenges by Gustavo Demarco and Rafael Rofman
9816	Getting an Earful: A Review of Beneficiary Assessments of Social Funds by Daniel Owen and Julie Van Domelen
9815	This paper has been revised, see Discussion Paper No. 9923
9814	Family Allowances by Suzanne Roddis and Zafiris Tzannatos

No.	Title
9813	Unemployment Benefits by Zafiris Tzannatos and Suzanne Roddis
9812	The Role of Choice in the Transition to a Funded Pension System by Robert Palacios and Edward Whitehouse
9811	An Alternative Technical Education System: A Case Study of Mexico by Kye Woo Lee
9810	Pension Reform in Britain by Edward Whitehouse
9809	Financing the Transition to Multipillar by Robert Holzmann
9808	Women and Labor Market Changes in the Global Economy: Growth Helps, Inequalities Hurt and Public Policy Matters by Zafiris Tzannatos
9807	The World Bank Approach to Pension Reform by Robert Holzmann
9806	Government Guarantees on Pension Fund Returns by George Pennacchi
9805	The Hungarian Pension System in Transition by Robert Palacios and Roberto Rocha
9804	Risks in Pensions and Annuities: Efficient Designs by Salvador Valdes-Prieto
9803	Building an Environment for Pension Reform in Developing Countries by Olivia S. Mitchell
9802	Export Processing Zones: A Review in Need of Update by Takayoshi Kusago and Zafiris Tzannatos
9801	World Bank Lending for Labor Markets: 1991 to 1996 by Amit Dar and Zafiris Tzannatos

Summary Findings

This paper presents the main approaches and tools to measuring poverty and vulnerability. This involves among other things the important issues of the setting of poverty lines, of the statistical use of surveys, of the empirical measurement of living standards, and of making poverty comparisons. The computation of aggregate poverty indices and the use of poverty dominance testing are also described. Throughout, implications for the understanding of the effects of public policy are highlighted, in particular those related to the design of safety nets.

HUMAN DEVELOPMENT NETWORK