
Demystifying Poverty Lines

BY JEAN OLSON LANJOUW

Poverty lines are widely perceived as occupying a central role in poverty analysis. In fact, setting a poverty line often receives the bulk of attention and intellectual effort in studies of poverty. This paper reviews the uses to which poverty lines are put. It discusses the different methods for constructing a poverty line and the implicit and explicit assumptions in each approach. The choice between various methods of construction is related to the intended use of the poverty line. The paper suggests that some of the attention paid to finding a “precise” location of the poverty line is misplaced. All poverty lines will retain an element of arbitrariness, and a convincing analysis of poverty is built on a whole sequence of steps with the poverty line being just one of them.

Introduction—Reasons for Setting a Poverty Line

Poverty alleviation efforts are best founded on a sound diagnosis of the underlying causes and dimensions of poverty. Across different countries, regions, communities or even families, the identity of the poor, the degree of their poverty, and its causes, will differ. In order to develop realistic policies for poverty alleviation in a given setting, it is essential to understand the nature of poverty in that specific setting.

A common component in virtually all approaches to poverty analysis is the setting of a poverty line. The most obvious purpose of a poverty line is to distinguish the poor from the nonpoor. This function as a threshold also has other applications.

Monitoring Poverty

A common reason for constructing a poverty line is to allow the calculation of poverty rates (for example, the proportion of the population that is poor, or some other more complex poverty measure). These poverty rates can then be used to make comparisons across groups and to monitor changes in poverty over time in order to inform policymaking. For example, comparisons of poverty rates for different regions within a country might help in the targeting of transfers or to determine the best locations for development expenditures. In addition, the success of poverty alleviation efforts could be

judged by tracking changes in poverty rates over time. The effect of other policies on the poor, such as liberalization or stabilization efforts, could also be assessed by looking at changes in poverty rates before and after implementation. For these comparisons to be meaningful, the poverty lines used in each setting being compared must represent the same welfare level.

Developing a Poverty Profile

Although useful for making comparisons, calculating poverty rates does not, by itself, bring one much closer to answering the more fundamental questions regarding the *determinants* of poverty. However, a poverty line can also be used to identify the poor as a group so that they can be focused on in greater detail. A poverty profile can be drawn up which describes the characteristics of those in poverty. This can then be used to investigate the causes of poverty. More immediately, the profile identifies correlates of poverty (such as location, ethnicity, occupational status and so on), which can be used by policy makers to reach the poor when detailed household-level information on income or expenditure is not available.

A Threshold for Entitlements

In some countries the poverty line has come to serve as a reference income or expenditure level to which state transfers are linked. Those with incomes below the poverty line are entitled to various publicly provided benefits. When it has this function, the fiscal implications of the location of the poverty line acquire considerable importance. It is clear that the higher the poverty line is set, the greater the public resources which will be needed to provide such targeted benefits. This role for a poverty line is controversial because it can lead to situations where the poverty threshold is determined, not so much by a notion of deprivation or need, as by the availability of government funds.

A Focus for Public Debate

A poverty line can play a political role by helping to maintain poverty as a focus of public attention. Atkinson (1993) contrasts the discussion of poverty in the United Kingdom with that in the United States, arguing that the existence of a fairly well-defined poverty line in the U.S. has helped to increase public discussion of the particular circumstances of the poor and how they have been changing over time.

Relative versus Absolute Poverty Lines

There are two main types of poverty lines. The first, a *relative* poverty line, is simply determined from a percentage cut-off point in the welfare distribution, such as the income or consumption level below which, say, 30 per cent of the population is located. Alternatively, it might refer to a cut-off point such as one-half the median income. This approach to setting the

poverty line is attractive in that it is both simple and transparent, and it is quite functional in terms of identifying a population sub-group upon which to focus attention.

There are two principal disadvantages to this approach, however. First, a relative poverty line is not terribly useful if one wants to monitor poverty over time or space. There is always a bottom 30 per cent of the population, even if living standards for the whole population have risen over time. Similarly, this approach does not allow for comparisons of poverty across regions. Second, the relative poverty line is essentially quite arbitrary. It is not clear why poverty should be defined in terms of one percentage point instead of another—and what percentage point is settled upon can have a bearing on the characteristics of the population subgroup designated as poor.

The second type of poverty line, an *absolute* poverty line, is explicitly linked to a specific welfare level. Anchoring the poverty line in this way allows one to make comparisons over time or across groups. Most countries that have officially recognized poverty lines define these in an absolute sense, interpreting them as a fixed standard of living.¹ Particularly in developing countries, the notion of some level of living below which one can be considered poor in an absolute sense (as opposed to being disadvantaged relative to the rest of society), would seem to have some empirical relevance. For these reasons, the discussion of the construction of poverty lines in this paper concentrates on an absolute interpretation. For more discussion of relative versus absolute poverty lines, and also the distinction between objective, external assessments of welfare versus subjective assessments, see Atkinson (1989), Atkinson and Micklewright (1992), Ravallion (1994) and Sen (1977).

Typically, an *absolute* poverty line will be expected to fulfil more than one of the four purposes listed above, and quite possibly all of them. This can give rise to some tension since the different uses to which the poverty line will be put can have different implications for how it should best be constructed. On the one hand, when the calculation of poverty rates is the primary goal, it is useful if the poverty line is set in an objective and scientific manner so as to establish the validity of the specific poverty rates that are calculated, based on it. On the other hand, when a poverty line is intended only to function as a frontier between the poor and the nonpoor so as to construct a profile, it may remain a bit “fuzzy.” This is because, unlike rates, poverty profiles do not tend to vary markedly as a result of moderate changes in the location of the poverty line. In order to contribute to the development of a poverty profile, it is sufficient that the general location of the poverty line commands broad support, and that the line accords, at least approximately, with conventional intuition. Similarly, while the extent to which poverty is a subject of popular debate depends on many factors aside from where the poverty line happens to be located, a poverty line which is clearly understood and which is easy to interpret, by laymen as well as experts, can help to encourage such debate. These latter purposes would suggest, therefore, that emphasis should be on intuition and simplicity.

The task of setting poverty lines often receives an inordinate amount of attention in poverty studies. This is probably because of the perceived importance of seemingly precise quantitative poverty rates. However, in recent years it has been increasingly recognized that poverty lines unavoidably retain an element of arbitrariness, and further, that poverty measures inevitably embody some implicit or explicit normative judgements which need not appeal to everyone.² Paradoxically, going to extreme lengths to establish the “objective” validity of a poverty line in order to yield “precise” poverty rates can end up undermining the success of the poverty line in fulfilling its other purposes. If, in order to satisfy the criterion of objectivity, the process followed in developing a poverty line becomes excessively complex, non-specialists may find it difficult to interpret and accept, and this may prevent it from contributing to the broader debate.

This paper argues that constructing poverty lines involves something of a balancing act: in order to allow convincing calculations of poverty rates, the lines should be perceived as objective and scientific; but the line should also be reasonably transparent and interpretable to a layperson in order to fulfil its other purposes. In practice, poverty lines occupy different positions along these various dimensions. Those which have enjoyed long histories and widespread use, such as the lines in use in India and the United States, have been rather simple in concept but are at least notionally anchored to objective criteria.

Outline of the Paper

In developing an absolute poverty line, welfare is assumed to be linked to the consumption of goods (and services). The basic idea in setting an absolute poverty line is to identify a basket of minimum essential consumption items. Those who do not have sufficient resources to obtain the basket are considered poor and those who do have sufficient resources are considered nonpoor. There are various approaches to setting an absolute poverty line. Mostly, these differ in how they decide which goods to include in the basket. This decision is typically broken down into stages. First, a minimum food basket is chosen to obtain a food poverty line. Then, an amount is added to the food poverty line to allow for essential nonfood consumption. These choices are discussed in the second and third sections. In the fourth section we discuss important issues in the construction of poverty lines which arise when poverty is being compared across regions, population groups or time. The fifth section briefly describes a method for testing the robustness of poverty comparisons, which to some extent reduces the need to specify a particular poverty line. The last section summarizes the main methods used, their strengths and weaknesses, and important points to keep in mind when considering trends in poverty rates.

The Food Poverty Line

There are two common approaches to setting the food component of the poverty line: least-cost and expenditure-based. Although we shall discuss them separately, in practice many poverty lines are a blending of both. Both approaches are linked to some decision about what are the minimum nutritional requirements per day for a reference person or household—a moderately active middle-aged male or a family of four, for example. (We return to the question of what level these minimum requirements should be and whether they should be the same for all persons in the fourth section.) In practice, minimums are typically defined in terms of energy requirements (represented by caloric intakes) and it is assumed that obtaining sufficient calories implies that other nutritional requirements are met.

Least-Cost Food Poverty Line

The least-cost food poverty line is obtained by selecting baskets of food items which are plausibly consumed in a given setting (i.e., not wildly inconsistent with prevailing tastes) and then calculating which basket yields the specified caloric minimum at the lowest cost, given prevailing prices. The cost of this basket defines the food poverty line. An example of this approach being used to set a food poverty line is in the World Bank (1995a) poverty assessment for the Kyrgyz Republic. Here, however, the lowest cost food basket was determined to meet not only caloric minimums, but also those for protein, carbohydrates and fats.

An attraction of the least-cost approach is that it does not require detailed data on household consumption, beyond some informal sense of what items could be reasonably included in the food basket. One only needs to know the prices for food items and their calorie content.³ However, there are also significant drawbacks to following this approach. Most important, it must be recognized that people having a level of food expenditure equal to the food poverty line would probably not, in fact, be consuming the minimum number of calories stipulated. People have strong preferences regarding food, and will not necessarily purchase the cheapest calories available. A least-cost food basket will almost certainly not accord with any person's actual eating habits. A second drawback is that the process of determining the least-cost basket can become a complicated linear programming exercise. This is particularly true if, in addition to setting a calorie threshold, one tries to explicitly incorporate micronutrient and other minimums as well. Not only does this make the setting of a poverty line complicated, it also makes the result non-transparent to the layman. Thus, whether the poverty line which resulted from the rather complicated linear programming exercise for the Kyrgyz Republic will enter into widespread public use, is still to be seen.

Expenditure-Based Food Poverty Line

Rather than beginning with the costs of various food items, the first step in specifying an expenditure-based food poverty line is to examine the actual food consumption patterns of some segment of the population. The foods consumed by this group are included in the basket, weighted by expenditure shares, and the quantities are then set so as to reach the minimum calorie level.

A poverty study using data from the 1995–96 Nepal Living Standards Survey (NLSS) provides an example. Columns (1) and (2) of Table 1 together indicate the average per capita daily consumption, in grams, of 37 main food items, looking only at households in the second to fifth deciles (the bottom of the distribution being discarded to avoid possible data errors). The third column indicates the calorie content per gram of each of the 37 items and the final column, which is the product of columns $(1 \times 2) \times 3$, gives the average number of calories obtained from each source. The sum of column (4) shows an average per capita calorie consumption of 1851 kcal. Suppose that a caloric minimum had been set at 2000 kcal. Then, to set the food basket, each of the quantities in column (1) would be multiplied by $(2000/1851)$ or 1.08.⁴ As in the least-cost approach, once the basket is chosen the goods are priced to obtain the food poverty line.

Looking only at the consumption of households in the lower deciles of the expenditure distribution ensures that expensive, luxury food items are not heavily represented in the basket. And by basing the composition of the basket on existing consumption patterns, the combination of food items included in the basket is clearly consistent with local tastes. The important implication of this feature is that people with food expenditure at the food poverty line are likely, in fact, to be consuming at least the minimum number of calories. Further, compared to a least-cost basket which identifies the lowest cost manner of obtaining a single nutritional component (i.e., calories), it is more reasonable to assume that a basket chosen in this way to meet minimum caloric requirements will also be nutritionally balanced. Observed consumption patterns are likely to have evolved in such a way that they provide a balanced diet.⁵

One disadvantage of this approach is that it requires detailed household survey data on food consumption, which measures not only food expenditures but also quantities consumed. In addition, in most developing countries much of the food consumed by households, particularly in rural areas, is home-produced. This home-produced consumption must not be overlooked; it has to be priced and added to food purchases in the market.

The expenditure-based approach is the more commonly used method of constructing a food poverty line. The food poverty line that underlies the official poverty line in the United States is one of the early examples based on a household survey of consumption (Orshansky, 1963, 1965, and U.S. Bureau of the Census, 1991). Derived in the early 1960s, the U.S. food poverty line represents the value of a 'low-cost food plan' developed by the Department of Agriculture, consistent with the Department's 1955 household survey of food consumption.⁶

Table 1**Food Basket Composition for the Nepal Poverty Line**

Food Item	(1) Grams (or mls) Per Day	(2) Edible Share	(3) Calories Per Gram	(4) Total Kcal
1. Fine rice	29.76	1.00	3.49	103.88
2. Coarse rice	216.1	1.00	3.45	745.60
3. Beaten rice	4.362	1.00	3.46	15.092
4. Maize	61.28	1.00	3.42	209.56
5. Maize flour	35.23	1.00	3.42	120.50
6. Wheat flour	87.78	1.00	3.41	299.33
7. Millet	34.00	0.79	3.31	88.903
8. Black Pulse	2.081	1.00	3.47	7.2222
9. Masoor	8.161	1.00	3.43	27.991
10. Rahar	1.099	1.00	3.35	3.6811
11. Gram	0.653	1.00	3.60	2.3509
12. Eggs	0.549	1.00	1.73	0.9495
13. Milk	38.00	1.00	0.67	25.458
14. Baby milk	0.014	1.00	4.96	0.0686
15. Curd	1.348	1.00	0.60	0.8086
16. Ghee	1.363	1.00	9.00	12.270
17. Vegetable Oil	0.204	1.00	9.00	1.8332
18. Mustard	7.367	1.00	9.00	66.302
19. Potatoes	31.02	0.85	0.97	25.577
20. Onions	6.542	0.95	0.50	3.1073
21. Cauliflower	4.528	0.70	0.30	0.9508
22. Tomatoes	2.562	0.98	0.23	0.5774
23. Bananas	3.911	0.71	1.16	3.2214
24. Citrus fruit	0.772	0.67	0.48	0.2483
25. Mangoes	5.486	0.74	0.74	3.0043
26. Apples	0.333	0.90	0.59	0.1767
27. Pineapple	0.095	0.60	0.46	0.0263
28. Papaya	1.736	0.75	0.35	0.4557
29. Fish	1.882	0.78	0.97	1.4241
30. Mutton	1.798	1.00	1.94	3.4883
31. Buffalo	1.882	1.00	0.86	1.6184
32. Chicken	1.086	1.00	1.09	1.1835
33. Salt	13.15	1.00	0.00	0.0000
34. Sugar	3.926	1.00	3.98	15.625
35. Gur	0.978	1.00	3.98	3.8910
36. Sweets	1.496	1.00	3.19	4.7711
37. Tea	0.721	1.00	0.00	0.0000
				1851

Sources: NLSS Household Survey (1996); Gopalan, Rama Sastri and Balasubramanian (1976).

Nonfood Expenditure

While having resources sufficient to meet one's food needs is clearly essential in determining the threshold below which people are judged to be poor, few would consider this sufficient to define the poverty line. People who are able to afford their food requirements but who lack the resources needed to purchase basic clothing and shelter, for example, would be considered by most to be deprived in a basic sense. Although there may be general agreement that a minimum basket should include items in addition to food, such as clothing and shelter, there is much less agreement about precisely which items to include.

In some cases, a specific bundle of nonfood items may be specified, an approach that is analogous to the methods of constructing a food poverty line outlined in Section 2. More commonly, however, a minimum level of nonfood expenditure is accommodated by scaling up the food poverty line by some specified multiple. For example, the multiple might be 1.5, implying a final poverty line that is 50 per cent higher than the food poverty line.

Directly Choosing a Nonfood Basket

One approach to specifying essential nonfood expenditure is to simply choose directly what nonfood items should be included in the basket. These items are priced, and the total gives an amount for minimum nonfood expenditure. This total is then added to the food poverty line that has already been developed to yield a final poverty line. While attractively straightforward, this method is rarely used in practice. One country where it has been attempted is Uzbekistan. Since 1992, minimum income levels have been calculated by the Ministry of Labour by adding explicitly defined and priced nonfood items to a food poverty line (World Bank, 1994).

An advantage of this approach is that it is simple and does not require detailed information about household consumption. All that is required is information on the prices of the chosen goods. However, an important disadvantage is that it is likely to be viewed as paternalistic and arbitrary. There is no objective standard, analogous to the minimum nutritional requirements, to which a choice of specific nonfood items can be linked. This makes it more difficult to get consensus that a specific basket of nonfood items is essential for a minimal standard of living. Even if one were to agree that clothing should be included, for example, there usually exists a wide variety of clothing items, of varying quality and price, and no clear guidance for choosing among them.

Scaling Up the Food Poverty Line

A second approach seeks to ground the nonfood component of the poverty line in observed consumption behaviour. This method avoids choosing directly the specific items that should be included in minimum nonfood expenditure. Instead, the food poverty line is simply scaled up by some factor to allow for the purchase of some essential nonfood items to reach a

final poverty line. There are two ways that this is generally done, both of which are based on observed consumption patterns.

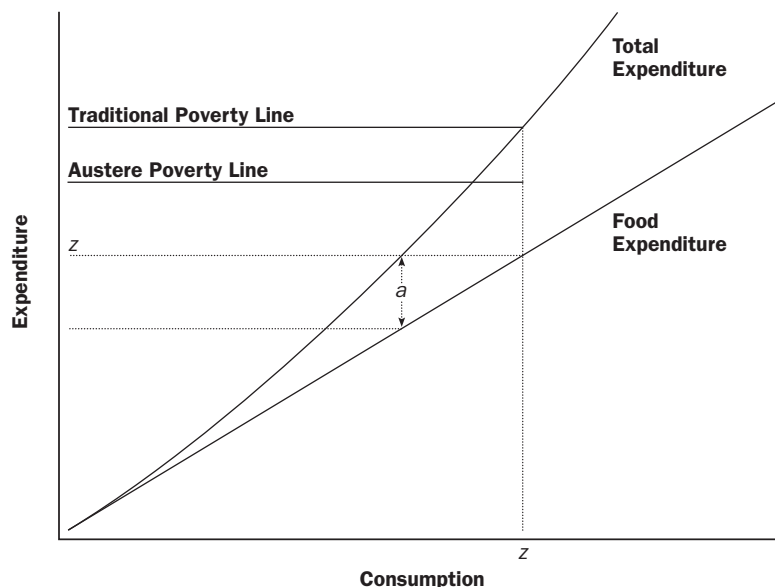
The most commonly used method is to determine the average level of total expenditure of those people whose food expenditures are just equal to the food poverty line. This level of total expenditure is then used as the final poverty line. (We shall call this the “traditional” poverty line.) The argument in favour of this method of reaching the final poverty line is that people with total expenditure below this level would be expected to have food expenditures below the food poverty line, and those with total expenditure above this level would be expected to have food expenditures above the food poverty line. (But there would be exceptions.)

An alternative, suggested by Ravallion (1994), is to determine the non-food expenditure of people whose *total* expenditure is equal to the *food* poverty line. This amount is then added to the food poverty line to obtain a final poverty line (which we shall call the “austere” poverty line). The argument made in favour of taking this approach is that if a household has the ability to obtain the minimum food basket but chooses to divert some of its resources to the purchase of nonfood items, then the household must clearly view those items as essential.

Households examined in the first case have total expenditure that is higher than the food poverty line, and so, higher than the households examined in the second case. Typically, then, these households would also have higher nonfood expenditures.⁷ As a result, the final poverty line obtained using the first method will be higher than that obtained using the second method. This is demonstrated in Figure 1. In the figure, the line marked *total* indicates the average total expenditure of households with any given level of food expenditure. The vertical distance between the lines marked *total* and *food* then represents nonfood expenditure. Suppose that the food poverty line is set at a value of z . To obtain the traditional poverty line we simply find the total expenditure of people spending z on food. For the austere poverty line we look for people whose total expenditure is z , measure their nonfood expenditure (the vertical distance $z-a$), and add it to z to obtain the final poverty line. The two final poverty lines are shown on the vertical axis.

Both methods of scaling up the food poverty line have some intuitive appeal and Ravallion (1995) suggests that they could be considered as upper and lower bounds for a reasonable final poverty line.

One practical question is how to find these final poverty lines if, as is likely, there is no group of people with total expenditure, or food expenditure, exactly equal to the food poverty line. One possibility is to estimate an econometric model of food expenditure as a function of total expenditure and other household characteristics. (This relationship between food and total spending is termed the *Engel curve*.) The resulting estimates may be used to predict the nonfood expenditure of households with a given level of food expenditure (see Ravallion, 1994, for a clear exposition).

Figure 1**Scaling Up the Food Poverty Line**

Another possibility is to use the following simple non-parametric procedure, used in obtaining a final poverty line for Nepal (see Lanjouw, Prennushi and Zaidi, 1997). The median per capita total expenditure was calculated for those households who had per-capita food expenditures in a small interval (plus or minus one per cent) around the food poverty line. Successively larger intervals were selected (a total of 5 times so that the largest interval was plus or minus 5 per cent) and then a simple average was taken of the five observations of median per-capita total expenditure around the food poverty line. This average then yields the traditional final poverty line.⁸

In closing this section, it is interesting to note that a recent study of poverty in Pakistan found that the two main alternative approaches to setting a final poverty line—one based on a full specification of both food and nonfood items, and the other following the more conventional methodology of scaling up the food poverty line—resulted in very similar final poverty lines (World Bank, 1995b). So, while these approaches are very different, they may yield quite similar results.

Comparisons of Poverty Rates

As discussed in the introduction, one of the most common reasons for setting a poverty line is in order to calculate poverty rates. These can then be used to monitor change over time, or differences across groups or regions. If this is one of the purposes of constructing a poverty line then welfare must be fixed across the groups being compared. This will involve scaling, that is, multiplying one or more group's welfare indicators by a constant so as to have them in the same units. A simple example is a comparison of

poverty at two points in time. If there has been inflation between the two periods then clearly it is not appropriate to use the same nominal poverty line for both periods.

Before proceeding, it is worth emphasizing that there are two ways that scaling can be done. Either the poverty line can be adjusted or, equivalently, the poverty line can be held constant and the welfare indicators, say incomes, which are being compared to the poverty line can be adjusted. For example, suppose we have a poverty line for 1997 at one dollar per person per day. Suppose that by 1998 inflation has caused prices on average to double. In comparing the poverty rates between 1997 and 1998, we must recognize that one dollar in 1998 can only buy half as many goods as that dollar in 1997. So, the poverty line for 1998 has to be increased in line with inflation to US\$2 per person per day. Alternatively, we could scale *down* the incomes in 1998 by one-half and keep the poverty line constant at US\$1 in both years. Of course, one has to be careful not to over-compensate by doing both. While scaling up the poverty line, or scaling down incomes are equivalent in theory, in practice there are conventions that have emerged, which depend on whether one is adjusting for inflation, demographic structure, spatial variation in prices, and so on. We will touch on this further below.

Price Adjustments and Consistency of the Basket

Poverty comparisons across different regions or between urban and rural areas are likely to require adjusting for variations in the cost of living. However, while most countries do publish consumer price indices that permit the temporal adjustment for cost of living variation, it is much less common to find similar price indices that reflect the fact that, at any point in time, prices can vary spatially to a considerable degree. Where poverty is being analyzed on the basis of a household consumption survey, it is often possible to construct a regional price index in order to achieve the spatial comparability of consumption or income measures.

A common issue that arises has to do with the fact that not only can prices vary markedly across regions, but consumption patterns can as well. When deriving the food poverty line, the conventional practice is to obtain some basket of goods, representing a certain nutritional value, which is consistent with the observed consumption patterns among low income households in the country *as a whole*. This common basket can then be priced using region-specific average prices per food item yielding region-specific food poverty lines.⁹ For reasons of presentational convenience, the convention is to construct a regional-cost of living index, and to adjust incomes or consumption expenditures by this index, and to apply these price-adjusted “real” incomes against a single, national, poverty line. In order to construct the regional cost of living index, the standard procedure is to calculate a national food poverty line as a population weighted average of the region-specific food poverty lines and then take the ratio of each regional food poverty line to the national food poverty line.

Allowing both food baskets, as well as prices, to differ across regions, or across the rural/urban divide, while appealing from a certain perspective, is not really acceptable because it makes it difficult to argue that the welfare level in the different regions is being held constant. The argument is sometimes made that insisting on a common consumption basket is unreasonable because consumption patterns across regions are very different. The critical question is then whether one believes that the reason for the divergence of consumption habits across, say, urban and rural areas is the result of taste differences or different levels of wealth in urban versus rural areas. For example, do consumers in urban areas consume more refined flour because they have fundamentally different taste patterns than rural households, or is it because they are, on average, wealthier and can more easily afford refined flour? If the latter is more likely, then a single consumption basket based on the consumption pattern of *low-income* households in the country *as a whole* may well be justified. It is quite possible that the higher average wealth in urban areas results in consumption patterns for urban areas which, on average, are quite different from those in rural areas, while at the same time, low-income households in urban areas might be consuming a basket of goods which is not all that different from that consumed in rural areas. In that case, taking the low-income average consumption basket and holding it constant would be quite defensible. On the other hand, if in one part of the country the staple diet of low income households consists mainly of potatoes, while in another the poor tend to eat bread, then one can imagine that taking a national average consumption basket would result in a mixture of potatoes and bread which, in fact, is not observed anywhere in the country. In this case the problem of different consumption patterns is less easy to circumvent. Even here, however, it is not clear that calculating different food poverty lines based on different consumption baskets is justified. (For further discussion of these issues see Bidani and Ravallion, 1994.)

Consumption surveys do not always include information on prices. Where they do not, it might still be possible to proceed with spatial cost-of-living adjustments by means of the unit values (expenditure per food item divided by quantity purchased) which can often be calculated from household surveys. These unit values are not the same as prices; it is difficult to distinguish actual price variation from quality differences (see Deaton, 1994). However, adjustments based on unit values are likely to remain more appealing than failure to adjust for cost of living variation altogether.

In principle, one could, and should, use the same approach in comparisons across time. This would involve taking the same food basket and repricing it in each period (and then recalculating the nonfood component based on these food poverty lines). However, in practice it is more common to simply multiply the final poverty line by some, already existing, cost of living index (such as the consumer price index or GDP deflator). This may not be appropriate because most general cost of living indices, such as the

consumer price index, are not based on the specific consumption patterns of the poor, and therefore do not capture well the price changes that are particularly relevant to them.¹⁰

Equivalence Scales

To monitor poverty over time and space, one must make adjustments to ensure that the basis for comparison is valid. Likewise, in order to compare different households one must make adjustments for different sizes and composition of households. During the discussion so far, the construction of a poverty line, particularly the food poverty line, has been linked to nutritional minimums for a reference person or household. For example, in the Nepal example discussed earlier, observed average household consumption of various food items was divided through by household size in order to obtain a measure of average per capita consumption of each food item per day. It was also assumed that the same per capita calorie requirement applied to all individuals, irrespective of age or gender. However, the welfare associated with a given basket of goods and services could well differ according to the size and demographic structure of the household.

As with cost-of-living adjustments, accounting for differences in household structure can be done either by constructing a separate poverty line for each type of household, or by scaling household income or consumption. For example, if we believe that a household comprising two adults and three children has different requirements than a household comprising only two adults, we could construct a different poverty line for these two types of households. Alternatively, we could adjust household income or consumption for these two types of households so as to make them comparable. While these two types of adjustments achieve the same objective, it is more common to find adjustments made to consumption or income rather than for different poverty lines to be specified for different types of households.

While it certainly seems likely that family members do not all require the same share of the household's total resources in order to reach the same welfare level, and that household resources are not allocated equally across all household members, there is no widely accepted alternative to the simple per capita convention. There is extensive literature that proposes a variety of alternative *equivalence scales* (rules for allocating household expenditure to household members of different gender and ages) but there exists little guidance for choosing among them.¹¹ A specification of the different "needs" of different family members, which is what equivalence scales attempt to summarize, can be based on (at least) two different methodologies: fixing the nutritional requirements of different types of people, or examining household consumption behaviour. Both of these have their drawbacks.

In developing countries it is common to find equivalence scales in use which are based on the different nutritional requirements of persons of different ages and gender. It might, for example, be the case that a child below the age of five is deemed to require only about one-third of the calories of

an adult male in order to be able to function normally. An equivalence scale based on nutritional norms might therefore suggest that a young child count as one third of an adult. In this way, a household comprising one adult male and three children below the age of five would be considered to comprise two “equivalent adult males.” There are extensive debates in terms of taking into account activity levels, the degree of physical adaptation of persons to different nutritional intakes at given activity levels, the importance of nutrients other than calories, and so on, in determining equivalence scales. Even if widespread agreement could be reached as to the precise nutritional requirements of different people it is certainly not obvious that a rule which applies to food intakes should be taken to apply to household expenditures as a whole. It is quite conceivable, for example, that while a child requires less in food than an adult, it requires more in terms of other components of household consumption such as education expenses, clothing and medicines. The application of nutrition-based equivalence scales to any expenditures other than food expenditures is thus essentially arbitrary. In a context where the consumption aggregate is fairly comprehensive, and food consumption is not an overwhelming component of total expenditure, this procedure is particularly unappealing.

The second approach to setting equivalence scales is based on empirical studies of household consumption behaviour. The fact that households have different consumption patterns when their composition differs is interpreted to reflect the different needs of persons of different ages and gender, and equivalence scales are then developed which summarize those needs (see Deaton and Muellbauer, 1986, for a clear exposition of this approach). As Atkinson (1989) points out, however, it is not clear whether one would want to interpret observed allocations of consumption expenditures to certain persons as reflecting their particular needs; the consumption patterns could also be reflecting patterns of discrimination (for reasons quite unrelated to actual requirements).¹² Equivalence scales based on observed consumption are therefore also controversial.

Choosing an alternative to the equal-share rule implicit in the per capita consumption calculations is thus problematic. Nevertheless, poverty rates tend to be quite sensitive to the introduction of equivalence scales. It is therefore advisable to calculate poverty rates several times using different equivalence scale, to see how much they change.

By contrast, it is quite often found that poverty profiles are much less sensitive. In a recent study for Ecuador, Hentschel and Lanjouw (1995) focus on the impact of introducing equivalence scales on the poverty profile. Table 2 describes what happens to a simple poverty profile for Ecuador, based on a 1994 household survey. In this table, the incidence of poverty is held constant at 35 per cent of the population, and the focus of attention is on the characteristics of the population that falls in this bottom segment of the consumption distribution. For example, in column (1) we see that

someone living in a rural area has a 47 per cent probability of being poor, someone living in a house with mud walls has a 43 per cent probability of being poor, and so on. Each column contains the poverty profile constructed with a different equivalence scale, as indicated in the column heading. A striking observation is that the poverty profile changes very little—even though the adjustments associated with the different equivalence scales are not minor. While some changes in the incidence of poverty associated with a particular household characteristic do occur, virtually all of the conclusions reached on the basis of an analysis of per capita consumption (column 1) carry through when any of the three alternative equivalence scales are applied. For the Ecuador data, one can therefore conclude that, while a decision to work with per capita consumption might have an influence on the calculated poverty rate, broad conclusions relating to the poverty profile are robust.¹³

Compatibility of Data

The issues discussed above are becoming increasingly well recognized. Purchasing power parity adjustments have been introduced to correct for varying costs of living across countries, and within countries. Similarly, in order to compare poverty over time, temporal price indices have been developed. Sensitivity analysis with alternative equivalence scales is now fairly common.¹⁴ Surprisingly, however, less attention tends to be paid to the importance of ensuring that the underlying welfare indicators being compared are defined in the same way.

When the poverty line is based on a different definition of well being than the available welfare measure at the household or individual level, then problems of comparability can arise. For example, suppose that the poverty line is defined in terms of household consumption expenditure. In one period the poverty rate is calculated using income as a welfare indicator, and in another period the calculation is carried out for the same poverty line but with household consumption as the welfare indicator. Any difference in measured poverty between the two periods could be entirely attributable to the fact that two different definitions of household welfare were used (and compared against the same poverty line).

A similar problem may arise even when only, say, consumption information is used but there are different degrees of detail available for the two data sets. Suppose, for example, that we wish to compare poverty between two countries where, for the first, we have total expenditure as an indicator of household welfare and, for the second, we have only information on food expenditure. Clearly, if poverty rates for the two countries are calculated with reference to the *same* poverty line, more people will tend to appear poor in the second country simply because fewer components of expenditure are included in the household welfare measure. One solution to this problem is to avoid a constant poverty line in the first place and calculate poverty

Table 2**The Effect of Alternative Equivalence Scales**

Household Characteristics	Base Per Capita Expenditure	Child = 0.5A	Infant = 0.3A Child = 0.5A	Infant = 0.3 Child = 0.5A Elderly = 0.5A
<i>Overall Poverty Incidence</i>	0.35	0.35	0.35	0.35
Rural	0.47	0.48	0.46	0.46
Urban	0.25	0.25	0.26	0.26
Sierra	0.33	0.33	0.33	0.33
Costa	0.35	0.35	0.35	0.35
Oriente	0.60	0.57	0.57	0.56
Black and White TV	0.34	0.35	0.34	0.35
Color TV	0.10	0.11	0.12	0.12
Bicycle	0.27	0.26	0.27	0.28
Refrigerator	0.14	0.16	0.17	0.17
Radio / Cassette	0.28	0.29	0.29	0.29
House with Mud Walls	0.43	0.43	0.42	0.41
House with Dirt Floor	0.59	0.59	0.58	0.58
House with Wood Walls	0.52	0.53	0.53	0.51
Telephone Connection	0.05	0.05	0.05	0.05
Networked Electricity	0.31	0.31	0.32	0.32
Networked Water	0.22	0.23	0.23	0.24
Waste Disposal	0.20	0.22	0.22	0.22
Sewage Removal	0.24	0.25	0.25	0.26
Head with No Education	0.58	0.58	0.59	0.57
Primary Educated Head	0.41	0.41	0.42	0.42
Secondary Educated Head	0.18	0.17	0.16	0.17
Tertiary Educated Head	0.05	0.04	0.04	0.05
Indigenous Head	0.68	0.70	0.68	0.68
Average Household Size				
Poor	5.84	5.56	5.52	5.67
Non-Poor	4.33	4.42	4.43	4.38

Note: A = adult.

Source: Hentschel and Lanjouw (1995).

lines for each country which are a function of the type of data available. But this must be done carefully. For example, returning to Figure 1, if only food expenditure information is available and the food poverty line is at the value z , then the number of people poor would be those with food expenditure less than z .

This would still be true if total expenditure information were available, *and* if the poverty line is scaled up in the traditional manner. Having total expenditure less than the traditional line implies that food expenditure is less than z and vice versa. Thus, the headcount rate of poverty is not affected by the comprehensiveness of the consumption measure as long as the poverty line is scaled up to reflect the aggregation of the data using the traditional approach. By contrast, it can be seen in the figure that this is not the case if scaling is done using the austere approach. Some households which have total expenditure above the austere line, and which would therefore be considered nonpoor if total expenditure were available as the welfare indicator, have food expenditure which is below the food poverty line z , and so would be considered poor if food expenditure alone were the welfare indicator. (Regardless of how the poverty line is scaled up, poverty measures other than the headcount will not, in general, be comparable. See Lanjouw and Lanjouw, 1997.)

This issue is important because it is very common to find that different data sets do not allow the application of identical definitions. For example, in economies undergoing rapid change (such as the transition economies of Eastern Europe and the former Soviet Union), price liberalization and the increased allocation of goods and services through market transactions over time have resulted in more items available for inclusion in a consumption aggregate.

It is also important to take account of this issue because comparisons can be strongly affected by ignoring it. Table 3 gives an example of the extent to which comparisons of poverty can be biased by differences in the comprehensiveness of the welfare indicator. The table shows the headcount measure of poverty for Pakistan calculated using different consumption aggregates. Panel A scales up the final poverty line using the traditional approach and panel B the austere approach. Panel A indicates that, if the poverty line is set using the traditional approach, the differences between the headcount ratio is robust to alternative definitions of consumption (the differences are not statistically significant). In panel B, where the austere approach to scaling up the poverty line is used, the poverty measures change substantially depending on the type of information available. The headcount falls from 61 per cent when only food expenditures are used, to 37 per cent when the most complete consumption expenditure information is available. Even comparing the rows that include both food and nonfood spending (but differ in degree of detail) the decline in the incidence of poverty is substantial. Note that the consumption aggregate defined in the second row of both panels is one that would already be considered comprehensive by most observers.

Table 3

**The Incidence of Poverty
Using Alternative Consumption Aggregations in Pakistan**

Consumption Aggregate	Poverty Line	Poverty	Incidence
<i>Panel A: Scaling up using the traditional approach</i>			
Food spending	245.7 Rs.	0.61	(0.01)
Food spending plus basic nonfood spending	336.7	0.59	(0.01)
Food plus basic nonfood spending, including energy and education spending	382.7	0.60	(0.01)
Above with imputed water expenditures	425.7	0.59	(0.01)
Above with imputed value of housing services	488.7	0.60	(0.01)
<i>Panel B: Scaling up using the austere approach</i>			
Food spending	245.7	0.61	(0.02)
Food spending plus basic nonfood spending	305.8	0.52	(0.02)
Food plus basic nonfood spending, including energy and education spending	326.8	0.48	(0.02)
Above with imputed water expenditures	336.7	0.40	(0.02)
Above with imputed value of housing services	355.5	0.37	(0.02)

Note: Estimated standard errors are in parentheses, and take into account stratification and clustering in the surveys.
(See Howes and Lanjouw, 1996.)

Source: Lanjouw and Lanjouw (1997).

There are two ways to deal with the problem of comparing data of differing levels of comprehensiveness. The first is to restrict attention to the headcount ratio using final poverty lines constructed in the traditional manner. This allows one to make comparisons over time or across regions without adjusting the data. See the appendix for an example from El Salvador. The second is to calculate predicted values of total expenditure for the periods or countries where actual total expenditure is not available, in order to have consumption measures which are comparable (Lanjouw and Lanjouw, 1997).

Robust Poverty Comparisons

Earlier in this paper we referred to an alternative approach to comparisons of poverty which is enjoying increasing popularity and avoids the need to set a specific poverty line. Rather than calculating specific poverty rates and then comparing these across countries, sub-groups or time-periods, entire distributions of income or consumption are compared, using various graphical techniques. Effectively, this “dominance-analysis” approach allows one to make poverty comparisons over a wide range of poverty lines (and over a variety of possible poverty measures). Recognizing that any single poverty line (and any single poverty measure) inevitably embodies certain normative viewpoints and assumptions, and is therefore never guaranteed to satisfy all possible observers, the dominance approach is very useful in that it allows one to find a range of poverty lines and poverty measures over which poverty comparisons are robust.

An example of how this simple approach is implemented is provided in Figure 2. In this example, the cumulative distribution functions of urban and rural Ecuador are plotted in a graph. For a given per capita consumption level on the horizontal axis, reading off the vertical axis for one of the curves indicates the incidence of poverty which would result if a poverty line equal to that per capita consumption level had been selected. For example, a poverty line of 65,000 sucres would imply a headcount rate of about 30 per cent in urban areas and about 60 per cent in rural areas. If, as in this figure, one curve (that representing urban Ecuador) always lies below the other curve, without any point of intersection, then results from the theory of stochastic dominance can be invoked for a fairly powerful conclusion: poverty in rural Ecuador is higher than in urban Ecuador *irrespective* of where you draw the poverty line.¹⁵ And, this is true for any poverty measure you might choose to employ. While the precise amount by which rural poverty in Ecuador exceeded that for urban Ecuador would depend on the choice of a particular poverty measure, and would require the choice of a particular poverty line, the dominance result reflected in the fact that the two distribution functions do not intersect anywhere allows one to make a very robust statement about the ranking of poverty across these two sectors.

Underlying this dominance approach to the measurement of poverty is the notion that by far the greatest value of poverty comparisons lies in an ability to *rank* poverty across settings, rather than come up with particular

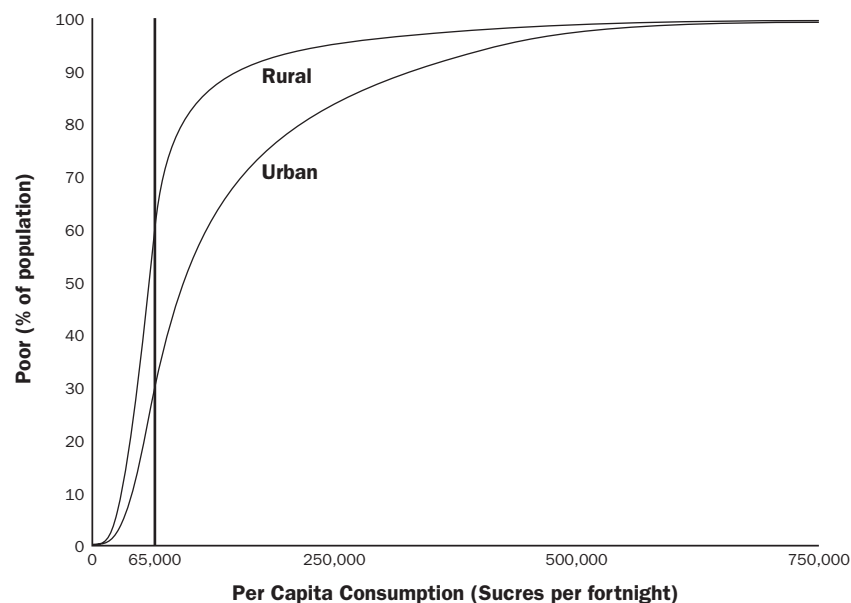
quantitative measures of the extent to which poverty across two settings differ. Cardinal comparisons are typically fragile to the specific assumptions that inevitably underlie a particular poverty line, or poverty measure. Simply examining distribution functions can, in principle, allow one to make much more robust comparisons of poverty and does not require one to settle on any one poverty line or measure.

An important caveat is that it is not always possible to reach such a broad conclusion as that shown in Figure 2. If the two curves were to intersect at some point, then it would no longer be the case that the same ranking of poverty would obtain over all possible poverty lines and measures. When distribution functions intersect, then one immediately knows that at least some rankings of poverty between the two sectors can be overturned by varying the choice of a particular poverty line or choosing an alternative poverty measure. Although there is no clear “dominance,” this conclusion also provides important information: it indicates that universal agreement may not be forthcoming regarding the ranking of poverty across the two sectors.

As the dominance analysis approach finds wider use, concomitantly less energy need be devoted to the derivation of a single poverty line. While this paper has emphasized other uses for a poverty line than simply an interest to calculate poverty rates (i.e., in order to construct a poverty profile, or in order to set public transfers) it is clear that the increased use of the dominance analysis approach to poverty measurement is likely to diminish the emphasis placed on the calculation of poverty lines.

Figure 2

Cumulative Distribution Functions in Urban versus Rural Ecuador (1995)



Summary

Poverty lines may be asked to fulfil a number of functions. Most commonly, the line is constructed in order for poverty rates to be calculated. To be convincing in this role, it is often thought that the poverty line should be scientifically and objectively derived. However, all poverty lines incorporate normative assumptions and elements of arbitrariness. No poverty line is therefore truly objective.

The other functions that a poverty line might be asked to include: identifying the poor for the construction of a poverty profile, creating a benchmark for public transfers or to stimulate public debate. Therefore, the line must be simple and easy to interpret. The most important property in this context is that the line receives widespread acceptance.

A successful poverty line must therefore balance the sometimes conflicting demands imposed by the line's differing roles. Too simple a line might appear arbitrary. But too much scientific precision can result in excessive complexity and a lack of transparency. In any event, the poverty line is only a first step in the larger programme of poverty analysis and policy formulation. This means that it should not receive excessive attention (and resources) at the expense of the arguably more important subsequent analysis.

There are numerous approaches to setting a poverty line. Typically these involve first determining a food poverty line and then arriving at an allowance for essential nonfood expenditures. The primary consideration in establishing the food poverty line is to what extent it will reflect actual consumption patterns of the poor and to what extent it will be weighted towards the lowest-cost calories (or some other nutritional criterion). At least some recognition of existing consumption patterns seems advisable.

Essential nonfood expenditure is sometimes determined by directly specifying a bundle of essential nonfood items. In practice it is difficult to make this choice in non-arbitrary manner. As a result, the more common approach is to simply scale up the food poverty line by some factor. This factor can be chosen in different ways from observed consumption patterns.

Table 4 summarizes the various steps and options in deriving a final poverty line. Additional issues arise when the purpose of setting poverty lines is to make comparisons across groups or over time. In this case it is essential that the level of welfare associated with the poverty line be the same in all settings which are being compared. It is important to correct for spatial and temporal variations in price using cost of living indices. Household size and composition should be accounted for using equivalence scales. Finally, it is essential to ensure that the underlying welfare definitions across settings are compatible.

New techniques of poverty analysis allow one to make comparisons of poverty without the need to define a specific poverty line or to settle on one particular poverty measure. To the extent that these techniques enter into widespread use, the emphasis and attention paid to the derivation of a poverty line is likely to diminish in the future.

Table 4
Summary of Approaches to Setting an Absolute Poverty Line

Component of the Final Poverty Line	Method	Advantages	Disadvantages
Food Poverty Line	Least-Cost	Identifies the lowest cost food bundle. Does not require detailed expenditure data.	Food bundle may not accord with actual eating habits. Complex—particularly with multiple nutritional minimums.
	Expenditure-Based	Consistent with eating habits of low-income households. More likely that obtaining a caloric minimum implies balanced nutrition.	Requires detailed household-level quantity and expenditure data.
Non-Food Component	Choosing Non-Food Bundle Directly	Straightforward and transparent. Does not require expenditure data.	Arbitrary and paternalistic. Requires price data.
	Scaling Up Food Poverty Line	Reflects behaviour of low-income households.	Requires household level expenditure data.

Appendix: Resolving Problems of Noncompatible Data

The following example demonstrates that comparisons of poverty rates are feasible even in the presence of concern about the comparability of the underlying consumption aggregates. If one is prepared to use the traditional approach to scaling up the poverty line, and to use only the headcount as a measure of poverty, then measured poverty rates will not be affected by the definitions of consumption used. In other words, when poverty rates do differ, this result will not have been driven by the definitions of consumption employed in the different data sets.

The *Encuesta de Hogares de Propósitos Múltiples* household survey in El Salvador (covering 4,229 households during the period July–September, 1994) was the subject of an experiment. Two non-overlapping samples were drawn from the same sampling frame and were administered using different consumption modules (both of which include some basic nonfood goods and services as well). A short module asking about the consumption of 18 food items and six nonfood items was completed for 3,182 households. A long module inquiring into the consumption of 72 food items and 25 nonfood items was completed for 1,047 households. Only with reference to five basic food items did the two coincide: corn tortilla, bread, sweet bread, beans and rice. Essentially, the long module referred, in a more detailed, itemized way, to broad food categories included in the short module.

Table 5 shows that average consumption levels by decile in the two sub-samples differ markedly, with households covered by the short module spending significantly less than those covered by the long module except in the top decile. If one were to measure poverty by applying a single poverty line across the two sub-samples, one would conclude that poverty among those covered by the short module is much greater than among those covered by the long module. Because both samples are representative of the same underlying population, this is clearly incorrect.¹⁶

Table 5

Per Capita Monthly Expenditure: Long and Short Questionnaires

Percentiles	Long	Short
10th	124.97	94.83
20th	193.60	153.35
30th	242.20	204.62
40th	296.78	255.24
50th	358.54	315.25
60th	445.25	382.19
70th	575.17	483.18
80th	730.99	627.97
90th	992.79	864.83
Top	2090.50	2225.40

Source: Republic of El Salvador (1994).

To deal with this problem, a food poverty line was defined based on only those food items about which exactly identical questions are asked in the two consumption modules and which are also included as components in the official (food) poverty line: tortilla maize, rice, beans and bread. These four items represent 32.5 per cent of the value of the official food poverty line, corresponding to 82.1 colones per person per month.

The final poverty lines derived from this (modified) food poverty line, using the traditional approach, are 575 colones per person per month for those households covered by the short consumption module, and 667 colones per person per month for those covered by the long consumption module. As expected, on this basis, 72 per cent of the population is poor in El Salvador irrespective of the consumption definition being used.

Footnotes

¹Absolute poverty lines can have a relative element if it is allowed that one's feeling of well being is, in part, a function of how one stands in relation to others in the same society. Indeed, one poverty line tradition (known as the "Leiden-School") explicitly canvasses the subjective views of households as to the level of income below which they would consider themselves as poor, in order to derive the poverty line (van Praag et. al., 1982, and Hagenaars, 1986).

²In this paper, attention will be confined mainly to the headcount measure of poverty. This is to simplify the exposition, not because this measure is preferred over alternative measures. The headcount is simply the percentage of the population below the poverty line. It is appealing because it is easily interpreted. However, it does have clear limitations. Most importantly the headcount index is entirely insensitive to the *degree* of poverty, that is, whether a person is just under the poverty line or far below it. In fact, a transfer of income from a very poor person to someone just below the poverty line could actually reduce the incidence of poverty even though the poorer person is worse off than before, if the recipient of the transfer crosses the poverty line as a result of the transfer. The *poverty gap* measure goes some way to address this shortcoming. This measure reflects the distance between the consumption levels of the poor and the poverty line, or the depth of poverty. A third measure, the *squared poverty gap*, differs in that it applies an increasing weight to distances below the poverty line. It is thus particularly sensitive to the severity of poverty. These three measures are all part of the Foster-Greer-Thorbecke (FGT) class of poverty measures. The poverty gap and squared poverty gap have not traditionally enjoyed the same widespread application as the headcount measure, probably because they are not as easily interpretable as the headcount measure. Nevertheless, their use is now becoming increasingly common. While these three are the most common poverty measures, there exists a wide variety of other candidates, each embodying a different set of assumptions about the importance of distance from the poverty line in measuring welfare (see Ravallion, 1994, for further discussion).

³Prices for the same good can vary depending on where they are purchased, and by whom. For this purpose, the relevant prices are those faced by relatively poor consumers.

⁴In fact, in addition to the 37 items listed, the NLSS solicited information on consumption of various "other," non-specified food items. It is not possible to calculate quantities consumed of these goods, although they contribute to total food consumption. This missing information was dealt with as follows: as shown in Table 1, consumption of the 37 items yielded 1851 kcals per person per day. It was also calculated that, on average, the 37 food items represented 85 per cent of total food spending. It was assumed that the calories obtained from "other" foods was proportionate to expenditure, that is, 15 per cent, yielding an average consumption level of $1851/0.85$ or 2178 kcals (see Lanjouw, Prennushi and Zaidi, 1997, for further details).

⁵This is not to say that attention need not be paid to other nutritional requirements—observed consumption patterns may not provide sufficient levels of some nutrients or may be imbalanced in other ways. For example, in the Kyrgyz Republic, where there is a strong tradition of eating meat, it was found that many energy-deficient households were consuming excessive fats and proteins even though the meat from which these derived was more expensive than calorie-intensive cereals.

⁶Deaton (1994) describes how the Orshansky poverty line has become widely accepted as a meaningful and empirically well-grounded poverty line. Yet he points out that this line is far from perfectly objective and non-arbitrary: the selected food plan itself was only one of several devised from the observed consumption patterns.

⁷This follows from Engel's Law, which states that the share of food in total consumption falls as total consumption expenditure increases.

⁸Following the austere approach, the median *food* expenditure would have been calculated for households with *total* expenditure falling within each of the intervals around the food poverty line.

⁹We have already mentioned that it is usually only food items, leading to a food poverty line, which are specified in quantity terms. The nonfood component leading to a final poverty line is usually simply some adjustment from the food poverty line. In this case, each regional food poverty line can be scaled up as described in section IIIb to obtain a final poverty line for each of the regions being compared. In those cases where physical quantities of nonfood items are also specified (as in the Uzbekistan poverty line) region-specific prices for these items should also be applied.

¹⁰India routinely calculates a separate cost of living index for agricultural labourers (the Consumer Price Index for Agricultural Labourers, CPIAL) which is more likely to capture changes in prices faced by the poor.

¹¹See for example, Buhmann, et al. (1988) and Coulter, et. al. (1992).

¹²Pollak and Wales (1979) have also demonstrated that the empirical models used to uncover these different needs are not well-identified; it is not possible to exclude very different interpretations for the observed consumption patterns.

¹³In a similar analysis for India, Dréze and Srinivasan (1997) find that the incidence of poverty among widow-headed households is not very sensitive to the specification of alternative equivalence scales related to the characteristics of family members. However, this is not the case when they consider the related issue of *economies of scale* in household consumption (that the per capita cost of reaching a particular level of welfare might be lower in large households than small households, through, for example, bulk-purchase discounts or the fact that many consumption items have public good attributes). When Dréze and Srinivasan adjust for the possibility of economies of scale they find that widow-headed households face a much higher risk of poverty than before the adjustment. This stems from the fact that widow-headed households are typically quite small and therefore benefit very little from economies of scale in consumption. (See also Lanjouw and Ravallion, 1995, and Hentschel and Lanjouw, 1995.)

¹⁴A recent study by a panel of experts on the United States poverty line (Citro and Michael, 1995) has proposed some modifications to the current U.S. poverty line. The panel's concerns include certain issues of the structure of the poverty line, notably regional cost-of-living variation, and adjustments for household size and composition.

¹⁵The theory of stochastic dominance was developed to offer guidance in comparing distributions which differ from each other in terms of both their variances as well as their mean values. For example, the theory can be used to compare portfolios of financial investments (which reflect different risk and expected return combinations) or alternatively to compare income distributions (which reflect different combinations of inequality and average income). There is now a large literature on the stochastic dominance approach to poverty and welfare comparisons. Useful references include Atkinson (1970, 1987, 1989), Deaton (1994), Ravallion (1994) and Shorrocks (1983).

¹⁶The two sub-samples were drawn from the same frame and were explicitly intended to be identical in all respects except for the consumption modules. Scott and Jolliffe (1995) show, in terms of location, household size, income levels, education, etc., that households in the two sub-samples resemble each other closely.

References

Atkinson, A. B., 'On the Measurement of Poverty', *Econometrica*. Vol. 55, 1987.

Atkinson, A. B., 'On the Measurement of Inequality', *Journal of Economic Theory*. Vol. 2, 1970, pp. 244–263.

Atkinson, A. B., *Poverty and Social Security*, Harvester Wheatsheaf, Hemel Hempstead, 1989.

Atkinson, A. B., 'The Institution of an Official Poverty Line and Economic Policy', Welfare State Programme Discussion Paper No. 98. STICERD, London School of Economics, 1993.

Atkinson, A. and F. Bourgignon, 'Income Distribution and Differences in Needs', in George R. Feiwel (ed.), *Arrow and the Foundations of the Theory of Economic Policy*, Macmillan Press, London, 1987.

Atkinson, A. B. and J. Micklewright, *Economic Transformation in Eastern Europe and the Distribution of Income*, Cambridge University Press, Cambridge, 1992.

Bidani, B. and M. Ravallion, 'How Robust is a Poverty Profile?' *World Bank Economic Review*, Vol. 8, 1994.

Buhmann, B., L. Rainwater, G. Schmaus and T. M. Smeeding, 'Equivalence Scales, Well-being, Inequality and Poverty', *Review of Income and Wealth*, Vol. 34, 1988.

Government of Nepal, *Summary Statistics of the Nepal Living Standards Measurement Survey*, Vol. II, Central Bureau of Statistics, 1997.

-
- Citro, C. and R. Michael, *Measuring Poverty: A New Approach*, National Research Council: National Academy Press, 1995.
- Coulter, F., F. Cowell and S. Jenkins, 'Equivalence Scale Relativities and the Extent of Inequality and Poverty', *Economic Journal*, Vol. 102, 1992.
- Deaton, A., 'The Analysis of Household Surveys: Micro-econometric Analysis for Development Policy', mimeo, Princeton University, 1994.
- Deaton, A. and J. Muellbauer, 'On Measuring Child Costs: With Applications to Poor Countries', *Journal of Political Economy*, Vol. 94, 1986, pp. 720–744.
- Dréze, J. and A. Sen, *Hunger and Public Action*, Oxford University Press, Oxford, 1989.
- Dréze, J. and P. V. Srinivasan (forthcoming), 'Widowhood and Poverty in Rural India: Some Inferences from Household Survey Data', *Journal of Development Economics*.
- Gopalan, C., B. V. Rama Sastri and S. C. Balasubramanian, *Nutritive Value of Indian Foods*, National Institute of Nutrition, Indian Council of Medical Research, Hyderabad, 1976.
- Hagenaars, A., *The Perception of Poverty*, North Holland, Amsterdam, 1986.
- Hentschel, J. and P. Lanjouw, 'Constructing an Indicator of Consumption for the Analysis of Poverty: Principles and Illustrations with Reference to Ecuador,' LSMS Working Paper, No. 124, Policy Research Department, World Bank, Washington, D.C., 1995.
- Howes, S. and J. O. Lanjouw, 'Making Poverty Comparisons Taking Into Account Survey Design: How and Why', LSMS Working Paper, World Bank, Washington, D.C., 1996.
- Lanjouw, J. O. and P. Lanjouw, 'Poverty Comparisons with Non-compatible Data: Theory and Illustrations,' Policy Research Working Paper No. 1709, World Bank, Washington, D.C., 1997.
- Lanjouw, P. and M. Ravallion, 'Poverty and Household Size,' *Economic Journal*, Vol. 105, No. 433, 1995.
- Lanjouw, P., G. Prennushi and S. Zaidi, 'Building Blocks for a Consumption-Based Analysis of Poverty in Nepal,' mimeo, World Bank, Washington, D.C., 1997.
- Marnie, S. and J. Micklewright, 'A Note on the Subsistence Minimum Income Levels Calculated by the Ministry of Labour, Uzbekistan,' mimeo, European University Institute, Florence, 1993.
- Orshansky, M., 'Children of the Poor', *Social Security Bulletin*, Vol. 26, 1963.
- Orshansky, M., 'Counting the Poor: Another Look at the Poverty Profile,' *Social Security Bulletin*, Vol. 28, 1965.

Pollak, R. A. and T. J. Wales, 'Welfare Comparisons and Equivalence Scales,' *American Economic Review*, Vol. 69, 1979.

Ravallion, M., *Poverty Comparisons*, Harwood, Chur, 1994.

Ravallion, M., 'Bounds for a Poverty Line,' mimeo, Policy Research Department, World Bank, Washington, D.C., 1995.

Republic of El Salvador, *Encuesta de Hogares de Propósitos Múltiples*, 1994-III, 1994.

Scott, K. and D. Jolliffe, 'A Comparison of Alternative Consumption Modules in the El Salvador Encuesta de Propósitos Múltiples', mimeo, PRDPH, World Bank, Washington, D.C., 1995.

Sen, A. K., 'On Weights and Measures,' *Econometrica*, Vol. 45, 1977.

Shorrocks, A., 'Ranking Income Distributions,' *Economica*, Vol. 50, 1983.

United States Bureau of the Census, *Poverty in the United States: 1991*, U.S. Government Printing Office, Washington, D.C., 1991.

van Praag, B., A. Hagenaars and J. van Weeren, 'Poverty in Europe,' *Review of Income and Wealth*, Vol. 28, 1982.

World Bank, *Uzbekistan: Adjusting Social Protection*, Report No. 13023-UZ, 1994.

World Bank, *The Kyrgyz Republic: Poverty Assessment and Strategy*, Report No. 14380-KG, 1995a.

World Bank, *Pakistan Poverty Assessment*, Report No. 14397-PAK, 1995b.