

The views expressed in this paper are those of the authors and do not necessarily reflect the views or policies of the Asian Development Bank. The Asian Development Bank does not guarantee the accuracy of the data presented.

Specifying Poverty Lines:

How and Why

BENU BIDANI, GAURAV DATT,
JEAN OLSON LANJOUW, and PETER LANJOUW

Benu Bidani is a researcher at the World Bank; Gaurav Datt is with the World Bank Sydney Office; Jean Olson Lanjouw is with Yale University; and Peter Lanjouw is with the World Bank. This paper is to be delivered at the ***Asia and Pacific Forum on Poverty: Reforming Policies and Institutions for Poverty Reduction***, to be held at the Asian Development Bank, Manila, 5-9 February 2001.

I. 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, below which persons are judged to be poor and above which they are not. In addition, a poverty line may also serve a number of other purposes.

A. 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 for making comparisons across different subgroups of the population and to monitor changes in poverty over time in order to inform policy making. 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.

B. 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 focussed on in greater detail. A poverty profile can be drawn up that 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 policymakers to reach the poor when detailed household-level information on income or expenditure is not available.

C. A Threshold for Entitlements

In some countries (notably in the United Kingdom and also increasingly in Eastern Europe and the countries of the former Soviet Union), 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 that 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.

D. 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 latter country has helped to increase public discussion of the particular circumstances of the poor and how they have been changing over time.

Section II discusses briefly the types of poverty lines: absolute, relative, and subjective poverty lines. For most purposes of public policy, absolute poverty lines are the most relevant, and the rest of the paper restricts its attention to such poverty lines. There are two main approaches to setting poverty lines, the first being the cost-of-basic needs method, which estimates the cost of a bundle of essential commodities. 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 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 alternative ways of setting such a 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 some amount is added to

the food poverty line to allow for essential nonfood consumption. These choices are discussed in section III.A (for the food poverty line), and section III.B (for the nonfood poverty line). Section IV discusses an alternate approach to setting poverty lines—the food energy intake method—which finds the expenditure level at which the household's food energy intake is just enough to meet the predetermined requirements. In section V we discuss important issues in the construction of poverty lines arising when poverty is compared across regions, population groups, or time. Section VI briefly describes a method for testing the robustness of poverty comparisons, which to some extent reduces the need to specify a particular poverty line. Section VII raises some other issues in setting poverty lines and section VIII concludes.

II. Types of Poverty Lines

A. Relative Poverty Lines

There are different types of poverty line. 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 percent of the population is located. Alternatively, it might refer to a cut-off point such as one half the median income, etc. 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 subgroup upon which to focus attention. A relative poverty line rises with the level of expenditure or income, hence richer regions have higher poverty lines.

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 percent 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 or countries. 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.


B. Absolute Poverty Lines

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, as it guarantees that two individuals with the same level of welfare would be treated the same way. For example, if two individuals had the same standard of living in all relevant aspects but are located in different regions, both would be treated as poor if their standard of living was lower than the poverty line for the region they lived in. Most countries with officially recognized poverty lines define these in an absolute sense, interpreting them as a fixed standard of living. 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.

C. Subjective Poverty Lines

A third less commonly used approach is to estimate subjective poverty lines. A strand of the literature has noted that there is an inherent subjectivity to the notion of basic needs, and has attempted to construct subjective poverty lines that have been based on the households' response to questions that elicit the minimum income needed to make ends meet (van Praag et al. 1982, Hagenaars 1986, Kapteyn et al. 1988). All households with incomes less than their reported minimum income are classified as poor. This clearly leads to inconsistencies, in that people with the same level of welfare may give different answers for their minimum income requirement, and hence may be treated differently, in terms of their classification into poor and nonpoor. In general, the empirical relationship between income and the minimum income reported by households is as shown in Figure 1. The subjective poverty line is selected as that income level at which people feel their income is just equal to the minimum income required. People with incomes above z^* feel that they have more than the minimum required, and people with incomes below this level feel they have less than the required amount.

Figure 1. **Subjective Poverty Line (z^*)**



Most of the literature attempting to use the minimum income requirements to derive subjective poverty lines has been done for developed countries. Such an approach has not been tried for a developing country. Measuring income in a developing country setting poses more serious challenges. An alternative approach has been tried using consumption adequacy in a recent paper by Pradhan and Ravallion (2000), which estimates subjective poverty lines for Jamaica and Nepal, using household survey responses by households on whether their food, housing, clothing, and total consumption were adequate for their families' need.

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).¹

For the rest of this paper, we will restrict our attention to absolute poverty lines. Typically, an absolute poverty line will be expected to fulfill 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 purposes for the poverty line have different implications for how it should best be constructed. For example, when the calculation of poverty rates is the primary goal, it is useful if the poverty line is set in an objective and

1. In this paper we indicate that an absolute poverty line represents a particular welfare level, and then express this in terms of income or consumption. It is of course clear that conventionally measured income or consumption are far from perfect proxies for individual welfare. While there exists much debate around this point (see, for example, Drèze and Sen 1989, and Atkinson and Bourguignon 1987) it suffices for our purposes to recognize that even these imperfect proxies are, in all likelihood, at least reasonably correlated with welfare.

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”. 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 that is clearly understood and 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 judgments that 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

2. In 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 alternatives 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. 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. They can be represented by the following expression:

$$P_{\alpha} = \sum_{i=1}^q \frac{(1 - x_i / z)^{\alpha}}{n}$$

where x_i is per capita consumption expenditure (or income) for those individuals below the poverty line, and zero for those above, z is the poverty line and n is the total population, and q is the number of poor people. P_{α} takes the value of zero for the headcount index, 1 for the poverty gap and 2 for the squared poverty gap. These measures are commonly used but there exist a wide variety of other measures, 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).

purposes. If, in order to satisfy the criterion of objectivity, the process followed in developing a poverty line becomes excessively complex, nonspecialists may find it difficult to interpret and therefore accept, and this may prevent it from contributing to the broader debate.

Constructing poverty lines thus 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 fulfill its other purposes. In practice, poverty lines occupy different positions along these various dimensions. Those that have enjoyed long histories and widespread use, such as the lines used in India and the United States, have been rather simple in concept but are at least notionally anchored to objective criteria.

III. Cost of Basic Needs Method for Constructing Poverty Lines

Poverty lines that are commonly estimated can be divided into two broad methods. The first method, called the Cost of Basic Needs (CBN) estimates the cost of a basket of essential consumption items. A second popular method, called the food energy intake (FEI) method, finds the expenditure (or income) level at which the household's food energy intake is just enough to meet the predetermined requirements. In Section IIIA, we discuss the derivation of the food poverty line in the cost-of-basic needs method, and in Section IIIB, we discuss the derivation of the nonfood part of the poverty line.

A. 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 section V.) 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.

(i) *Least-Cost Food Poverty Line*

The least-cost food poverty line is obtained by selecting baskets of food items that 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 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 minima, 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 importantly, 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 minima as well. Not only does this make the setting of a poverty line complicated, it also makes the result nontransparent to the layman. Thus, whether the poverty line that resulted from the rather complicated linear programming exercise for the Kyrgyz Republic will enter into widespread public use is still to be seen.

3. Prices 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.

(ii) *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 1995-1996 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)-(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 kcals. Suppose that a caloric minimum had been set at 2000 kcals. 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 that identifies the lowest cost manner of obtaining a single nutritional component (e.g., calories), it is more reasonable to assume that a basket chosen in this way to meet minimum caloric requirements will also be nutritionally

4. In fact, in addition to the 37 items listed, the NLSS solicited information on consumption of various "other", nonspecified 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 percent of total food spending. It was assumed that the calories obtained from "other" foods was proportionate to expenditure, that is, 15 percent, yielding an average consumption level of $1851/(.85)$ or 2178 kcals (see Lanjouw, Prennushi and Zaidi 1997, for further details).

balanced. Observed consumption patterns are likely to have evolved in such a way that they provide a balanced diet.⁵

But choosing the reference group is a crucial decision since the consumption pattern of this group is an important decision, as it determines the weights of the commodities in the consumption basket. The general approach is that researchers start with a prior belief about the level of poverty and use the consumption basket of that group to derive the food poverty line. An alternative approach used recently by Pradhan, Suryahadi, Sumarto, and Pritchett (2000) for Indonesia uses an iterative method to solve for the reference group. The method uses an initial reference group to estimate the food poverty line. The poverty line from this step is used as the center of the reference group for the next iteration. The method converges when the reference group gives a poverty line that is the same as the midpoint of the reference group. For the details of this methodology and the resulting reference consumption bundle, see Pradhan, Suryahadi, Sumarto, and Pritchett (2000).

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 US food poverty line represents the value of a “lowcost food plan” developed by the Department of Agriculture, consistent with the Department’s 1955 household survey of food consumption.⁶

5. This 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.

6. Deaton (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 nonarbitrary: the food plan itself was only one of several devised from the observed consumption patterns.

Table 1. Food Basket Composition for the Nepal Poverty Line

	Food Item (1)	Per Day Share (2)	Grams (or mls)		Total
			Edible Per Gram (3)	Calories Kcals (4)	
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	<u>0.0000</u>
					1851

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

B. 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 what items to include.

In some cases a specific bundle of nonfood items may be specified, an approach analogous to the methods of constructing a food poverty line outlined in section III. 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 percent higher than the food poverty line.

(i) 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 tried 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, it can also 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.

(ii) *Scaling Up the Food Poverty Line*

A second approach seeks to ground the nonfood component of the poverty line in observed consumption behavior. This method seeks to avoid 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 favor 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.

An alternative, suggested by Ravallion (1994), is to determine the nonfood 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 favor 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 2. 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

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

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.

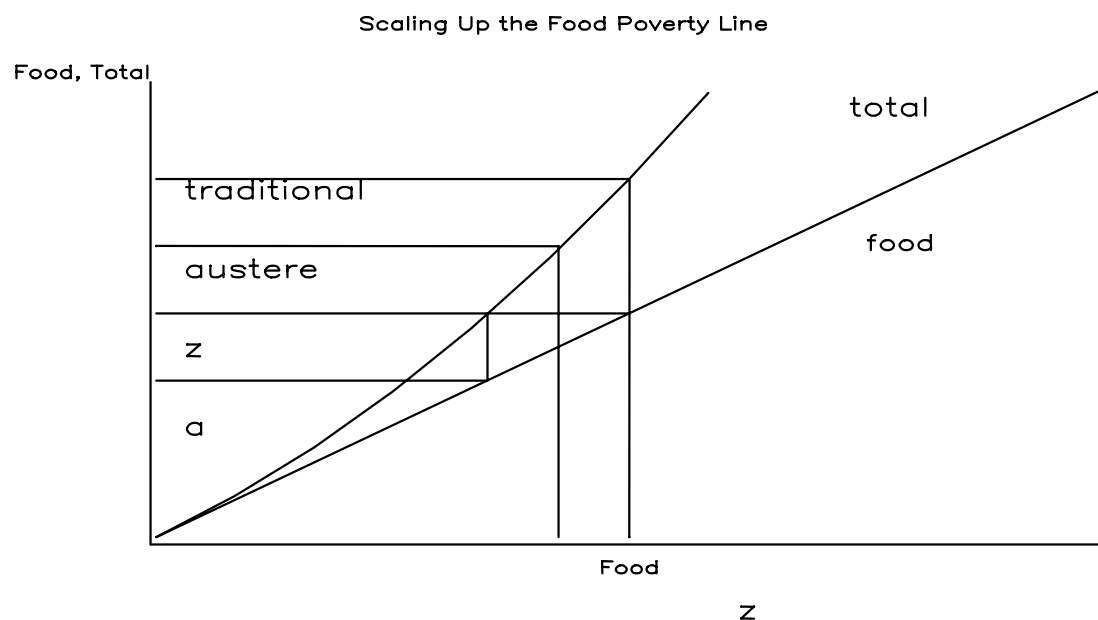
Both methods of scaling up the food poverty line have some intuitive appeal and Ravallion (1995) suggests that they could be considered as upper and lower bounds for a reasonable final poverty line.

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).

Another possibility is to use the following simple nonparametric 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 (of plus or minus one percent) around the food poverty line. Successively larger intervals were selected (a total of five times so that the largest interval was plus or minus 5 percent) 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

8. Following 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.

Figure 2




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.

IV. Food Energy Intake Method

The food energy intake method anchors the poverty line to the food energy requirements. This method estimates the total consumption or income at which the food energy intake is just enough to meet the predetermined food energy requirements. Setting the food energy intakes can be difficult (see section V). This method has two potential advantages: first, it does not need a separate estimate of the food and nonfood consumption needs as the method calculates the total expenditure needed for reaching the food energy intake. Second, the method does not require any information on prices. Overall, then the method has more modest data requirements.

Figure 3. The Food Energy Intake Method



But, the real question is, Does this method allow us to make comparisons across different subgroups of the population using a common yardstick for welfare? The answer is that this is unlikely for different reasons. The method has several problems. The relationship between the food energy intake and total consumption expenditure is likely to differ according to differences in tastes, activity levels, relative prices or publicly provided goods. And, there is nothing in the FEI method to guarantee these differences would be considered relevant to poverty comparisons (see Ravallion 1994, 1998 and Ravallion and Bidani 1994).

The following are examples that illustrate the particular problems associated with the food energy method. To the extent that prices differ between urban and rural areas (say due to transport costs), different nominal poverty lines should be used. But relative prices may also vary. For example, the prices of some nonfood goods tend to be lower in relation to food in urban areas than in rural areas, and retail outlets for nonfood goods also tend to be more accessible, so the full cost tends to be lower in urban areas. This may mean that the demand for food and hence food energy intake will be lower in urban areas than in rural areas at any given real expenditure level. But this does not, of course, mean that urban households are poorer at a given expenditure level.

Activities in typical urban jobs tend to require fewer calories to maintain body weight than do rural activities. So, food intakes will tend to be lower for urban workers at a given level of expenditure, but this should not be taken as a sign of poverty.

Tastes may also differ systematically. At given relative prices and real total expenditures, urban households may have more expensive food tastes; they eat more rice and less cassava and more animal protein and less food grains, or simply eat out more often. Thus, urban households pay more for each calorie, or their food energy intake will be lower at any given real expenditure. Again, it is not clear why we would designate a person who chooses to buy fewer and more expensive calories as poorer than another person at the same real expenditure level.

In each of the cases, the real expenditure level at which an urban resident typically attains any given caloric requirement will tend to be higher than in rural areas. And, this can hold true even if there is no cost of basic needs is the same across urban and rural areas.

An example from Indonesia (Ravallion and Bidani 1994) shows that there is little correlation in the regional rankings of poverty obtained by using FEI-based poverty lines and those which tried to hold constant purchasing power over basic consumption needs. And, the poverty estimates for urban and rural Indonesia from the two methods are diametrically opposed. The FEI-based poverty lines for urban and rural areas generate poverty rates for urban areas that exceed that for rural areas. The CBN method reveals that the headcount index for poverty is twice as high in rural areas as it is in urban areas.

V. Comparisons of Poverty Rates

As discussed in the introduction, one of the most common reasons for setting a poverty line is 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. 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 set in period one at z . By period two, the price level has doubled. In comparing the poverty rates in the two periods, one can either scale up the poverty line in the second period to $2z$, in line with the rate of inflation, or one may scale *down* the incomes

in the second period by one half and keep the poverty line constant at z . Of course, one has to be careful not to overcompensate by doing both. The reason that this is worth noting is that, while they are equivalent in theory, in practice the convention followed depends on the type of adjustment being made.

A. 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 so as 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

9. We have already mentioned that it is usually only food items, leading to a food poverty line that are specified in quantity terms. The nonfood component leading to a final poverty line is usually simply some scalar of 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.

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 that 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 potato, 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 potato 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 Ravallion and Bidani 1994).

The issue of one fixed national basket versus having region-specific food bundles has been a big issue in Indonesia. Recent estimates for Indonesia constructed by the Biro Pusat Statistik (BPS) uses separate urban and rural reference groups that are chosen under an assumption of higher cost of living in urban areas rather than rural areas. The reference group for urban areas is households with per capita expenditures of Rp. 80,000-100,000 per month and the reference group for rural households is Rp. 60,000-Rp. 80,000 per month. The resulting poverty line for urban areas is 27 percent higher than that for rural areas. Alternate poverty lines, using the iterative method (see section III) that uses a single national basket, lead to an urban poverty line 11 percent higher than

the rural poverty line (Pradhan, Suryahadi, Sumarto, and Pritchett 2000). As Table 2 shows, the resulting poverty estimates using the iterative method show a much larger difference in the poverty rates across urban and rural areas than does the BPS method.

Table 2. **Urban–Rural Differences in Iterative and BPS Methods for Indonesia, February 1999**

	Poverty Line (Rp/Month)	Headcount Index of Poverty (%)
Iterative method:		
Urban	90,490	16.34
Rural	81,184	34.10
Ratio (urban/rural)	1.11	0.48
BPS Method:		
Urban	93,869	19.98
Rural	73,898	25.85
Ratio (urban/rural)	1.27	0.77

Source: Pradhan, Suryahadi, Sumarto, and Pritchett (2000).

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. However, in practice it is more common to simply scale the poverty line by some, already existing, cost of living index. This is not appropriate because generic cost of living indices are typically based on the consumption patterns of the whole population. If anything, they tend to be weighted toward consumption patterns in wealthier, urban areas due to sample survey coverage.¹⁰

B. **Adjusting for Differences in Household Size and Composition:**

Equivalence Scales and Economies of Scale

During the discussion so far, the construction of a poverty line, particularly the food poverty line, has been linked to nutritional minima for a reference person or household. For example, in the

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

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 a large 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 between 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

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

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”. Even if widespread agreement could be reached as to the precise nutritional requirements of different people (and 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) it is certainly not obvious that a rule applying 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, she requires more in terms of other components of household consumption such as education expenses, clothing, medicines, and so on. The application of nutrition-based equivalence scales to any expenditure 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 behavior. 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 to summarize those needs (see Deaton and Muelbauer 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 so it is advisable to test how sensitive calculated rates are by using a number of different equivalence scales.

12. Pollak 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.

By contrast, it is quite often found that poverty profiles are much less sensitive. In a study for Ecuador, Hentschel and Lanjouw (1995) focus on the impact of introducing equivalence scales on the poverty profile. Table 3 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 percent of the population, and the focus of attention is on the characteristics of the population falling 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 percent probability of being poor, someone living in a house with mud walls has a 43 percent 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

C. Economies of Scale

The implicit assumption one makes when calculating per capita income or per capita expenditure is not only that the needs of different household members are the same but also that the cost of reaching a particular standard of living per person does not vary with the size of the household. The notion that there might exist economies of scale in household consumption, such that a certain standard of living is reached at a lower cost per person in a larger household than in a small household would appear more reasonable the larger the share of public goods purchased by

13. In a similar analysis for India, Drèze 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 Drèze 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.)

the household. Examples of such public goods might include housing services, consumer durables, and so on. In addition, the larger the size of households on average, the greater the scope for economies of scale to affect comparisons of welfare across households of different sizes. If there are economies of scale in consumption, then an analysis based on per capita consumption or income would tend to systematically overstate the poverty of larger households relative to small households. In developing countries the assumption of no, or negligible, economies of scale has typically been justified by pointing to the large share of food in total consumption and the notion that food is clearly a private good. On the other hand, however, it is also the case that family sizes in developing countries are on average relatively large¹⁴.

Empirically ascertaining the extent of economies of scale in consumption is difficult and prey to the same methodological problems as in empirically calculating equivalence scales. However, one attempt, for Pakistan, suggested that the extent of such economies may be considerable (Lanjouw and Ravallion 1995)¹⁵. Once again, the best course of action is probably to conduct sensitivity analysis - and to focus most closely on the impact of alternative scale parameters on the underlying poverty profile. Interestingly, when such sensitivity analysis was undertaken with data for Pakistan and Ecuador, the underlying poverty profiles proved, with one notable exception, to be quite robust. While a solid finding from many poverty studies based on per-capita measures of welfare is that large households tend to be poorer than small households, this finding was overturned for Pakistan or Ecuador once allowance was made for a fairly sizeable degree of scale economies. However, other demographically-linked elements of the poverty profile, such as that households with a larger dependency ratio are more likely to be poor was robust even in the presence of sizeable scale economies (Hentschel and Lanjouw 1995).

14. We have seen that in Kyrgyz Republic, the share of total consumption expenditure devoted to food is 71 percent, a comparatively high level even for developing countries. Average family size in Kyrgyz Republic is 4.7 members, which relative to other developing countries is not particularly large. These two factors might suggest that the scope for significant scale economies in consumption is rather limited. However, it is clear that with economic transition and possibly a rising share of total spending on nonfood items as more of such goods and services are transacted in the market, the scope for scale economies is likely to increase.

15. It is interesting to note that poverty lines calculated for households of different sizes following the subjective poverty measurement approach also suggest very sizable economies of scale (Atkinson 1991)

Table 3. 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
Overall Incidence Poverty	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
Nonpoor	4.33	4.42	4.43	4.38

Note: A = adult.

Source: Hentschel and Lanjouw (1995). **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.¹ 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 lines for each country that are a function of the type of data available. But this must be done carefully. For example, returning to Figure 2, 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 . If, on the other hand, total expenditure information were available, *and* if the poverty line is scaled up in the traditional manner, the same households would be identified as poor. 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 with 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 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 4 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 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 percent when only food expenditures are used to 37 percent 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 4. **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).

VI. Robust Poverty Comparisons

Earlier in this paper we referred to an alternative approach to comparisons of poverty that 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, subgroups 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.¹⁶

An example of how this simple approach is implemented is provided in Figure 4. 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 that would result if a poverty line equal to that per capita consumption level had been selected. For example, a poverty line of 50,000 sucres would imply a headcount rate of about 30 percent in urban areas and about 60 percent 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

16. There is now a large literature on this approach, which is applicable not only to poverty measurement but also to the measurement of inequality and welfare in general. Useful references include Atkinson (1970, 1987, 1989), Deaton (1994), Ravallion (1994), and Shorrocks (1983).

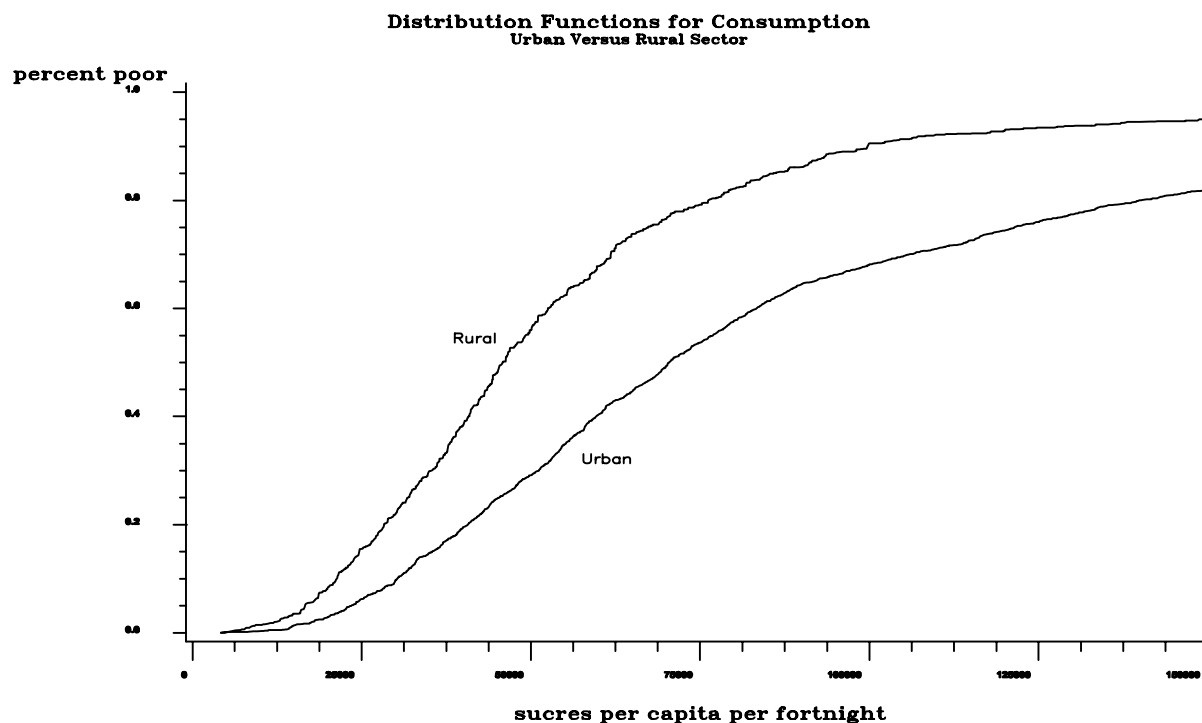
you draw the poverty line. And this is true for any poverty measure you might choose to employ. While the precise amount by which rural poverty in Ecuador exceeded that for urban Ecuador would depend on the choice of a particular poverty measure, and would require the choice of a particular poverty line, 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 might vary. 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 4. 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. The poverty comparison is not robust to choice of the poverty line or 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 purposes for a poverty line than simply an interest to calculate poverty rates (e.g., 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 4. Ecuador: Distribution Functions for Consumption



VII. Other Issues

A. Updating Poverty Lines over Time:

There is little consensus on the issue of how often poverty lines should change. That is, as countries develop and average incomes grow, should the real value of the poverty line also increase? Several countries, such as India, Indonesia, and US, all of which have been monitoring poverty over the past two decades, have been using constant real poverty lines (Ravallion 1998), though this practice has been questioned. But the process is a slow one, and there is little consensus on how often this adjustment should be made. Recent discussions in Indonesia on a strategy for poverty reduction has been accompanied by a dialogue on what poverty means today in Indonesia and a new method for calculating the poverty line has resulted in a higher real poverty line¹⁷.

17. Indonesia's Central Bureau of Statistics has revised its methodology for deriving the poverty line, which raised the headcount index of poverty in Indonesia from 11.3 percent in February 1996 (under the old method) to 17.7 percent under

VIII. Conclusion

Poverty lines may be asked to fulfill 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 fulfill: identifying the poor for the construction of a poverty profile, creating a benchmark for public transfers, or to stimulate public debate, emphasize that other characteristics of a poverty line may be equally important. In particular, they may require that the line 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 program of poverty analysis and policy formulation.

There are numerous approaches to setting a poverty line. Typically these involve determining first 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 toward 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 specifying directly a bundle of essential nonfood items. In practice it is difficult to make this choice in a nonarbitrary 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.

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 being compared. It is important to correct for spatial and

the new method. The World Bank has also estimated an alternate methodology using the consumption basket of households in 1999.

temporal variations in price using cost of living indices. 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 specify 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.

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 datasets.

The *Encuesta de Hogares de Propósitos Múltiples* household survey in El Salvador (covering 4220 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 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 enquiring into the consumption of 72 food items and 25 nonfood items was completed for 1047 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 A.1 shows that average consumption levels by decile in the two subsamples differ markedly, with households covered by the short module spending significantly less than those covered by the long module. 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.¹⁸

18. The 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.

Table A.1. **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% 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% of the population is poor in El Salvador irrespective of the consumption definition being used.

References

- Atkinson, A.B. (1987) 'On the Measurement of Poverty,' *Econometrica*. Vol. 55.
- Atkinson, A. B. (1970) 'On the Measurement of Inequality,' *Journal of Economic Theory*. Vol. 2, pp. 244-63.
- Atkinson, A.B. (1989) *Poverty and Social Security* (Hemel Hempstead: Harvester Wheatsheaf).
- Atkinson, A.B. (1993) 'The Institution of an Official Poverty Line and Economic Policy,' Welfare State Program Discussion Paper no. 98. STICERD, London School of Economics.
- Atkinson, A. and F. Bourguignon (1987) 'Income Distribution and Differences in Needs,' in George R. Feiwel (ed.) *Arrow and the Foundations of the Theory of Economic Policy*. London: Macmillan Press.
- Atkinson, A.B. and J. Micklewright (1992) *Economic Transformation in Eastern Europe and the Distribution of Income* (Cambridge: Cambridge University Press).
- Buhmann, B., L. Rainwater, G. Schmaus and T.M. Smeeding (1988) 'Equivalence Scales, Well-being, Inequality and Poverty,' *Review of Income and Wealth*. Vol. 34.
- Government of Nepal (1997) *Summary Statistics of the Nepal Living Standards Measurement Survey*. Vol. II. Central Bureau of Statistics.
- Citro, C. and R. Michael (1995) *Measuring Poverty: a New Approach* (National Research Council: National Academy Press).
- Coulter, F., F. Cowell and S. Jenkins (1992) 'Equivalence Scale Relativities and the Extent of Inequality and Poverty,' *Economic Journal*. Vol. 102.
- Deaton, A. (1994) 'The Analysis of Household Surveys: Microeconomic Analysis for Development Policy,' mimeo, Princeton University.
- Deaton, A. and J. Muellbauer (1986) 'On Measuring Child Costs: with Applications to Poor Countries', *Journal of Political Economy*. Vol. 94, pp. 720-44.
- Drèze, J. and A. Sen (1989) *Hunger and Public Action* (Oxford: Oxford University Press).
- Drèze, 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 (1976) *Nutritive Value of Indian Foods* (Hyderabad: National Institute of Nutrition, Indian Council of Medical Research).
- Hagenaars, A. (1986) *The Perception of Poverty* (Amsterdam: North Holland).
- Hentschel, J. and P. Lanjouw (1995) '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.
- Howes, S. and J.O. Lanjouw (1996) 'Making Poverty Comparisons Taking Into Account Survey Design: How and Why,' LSMS Working Paper. The World Bank.
- Kapteyn, A., P. Kooreman and R. Willemse (1988), 'Some Methodological Issues in the Implementation of Subjective Poverty Definitions', *The Journal of Human Resources*, 23: 222-242.
- Lanjouw, J.O. and P. Lanjouw (1997) 'Poverty Comparisons with Noncompatible Data: Theory and Illustrations,' Policy Research Working Paper no. 1709, The World Bank.
- Lanjouw, P. and M. Ravallion (1995) 'Poverty and Household Size,' *Economic Journal*. Vol. 105, no. 433.
- Lanjouw, P., G. Prenzushi, and S. Zaidi (1997) 'Building Blocks for a Consumption-Based Analysis of Poverty in Nepal,' mimeo, the World Bank.
- Marnie, S. and J. Micklewright (1993) 'A Note on the Subsistence Minimum Income Levels Calculated by the Ministry of Labour, Uzbekistan,' mimeo, European University Institute, Florence.
- Orshansky, M. (1963) 'Children of the Poor,' *Social Security Bulletin*. Vol. 26.
- Orshansky, M. (1965) 'Counting the Poor: Another Look at the Poverty Profile,' *Social Security Bulletin*. Vol. 28.
- Pollak, R.A. and T.J. Wales (1979) 'Welfare Comparisons and Equivalence Scales,' *American Economic Review*. Vol. 69.

-
- Pradhan, M. and M. Ravallion (2000) 'Measuring Poverty Using Qualitative Perceptions of Consumption Adequacy,' *The Review of Economics and Statistics*, 82(3): 462-471.
- Pradhan, M., A. Suryahadi, S. Sumarto and L. Pritchett (2000), 'Measurements of Poverty in Indonesia: 1996, 1999, and Beyond,' Policy Research Working Paper No. 2438, The World Bank, Washington D.C.
- Ravallion, M. (1994) *Poverty Comparisons* (Chur: Harwood).
- Ravallion, M. (1995) 'Bounds for a Poverty Line,' mimeo, Policy Research Department, World Bank.
- Ravallion, M. (1998) 'Poverty Lines in Theory and Practice', LSMS Working Paper No. 133, World Bank, Washington. D.C.
- Ravallion, M. and B. Bidani (1994) 'How Robust is a Poverty Profile?' *World Bank Economic Review*. Vol. 8.
- Republic of El Salvador (1994) *Encuesta de Hogares de Propositos Múltiples, 1994-III*.
- Scott, K. and D. Jolliffe (1995) 'A Comparison of Alternative Consumption Modules in the El Salvador Encuesta de Propositos Múltiples,' mimeo. PRDPH, World Bank.
- Sen, A.K. (1977) 'On Weights and Measures,' *Econometrica*. Vol. 45.
- Shorrocks, A. (1983) 'Ranking Income Distributions,' *Economica*. Vol. 50.
- United States Bureau of the Census (1991) *Poverty in the United States: 1991* (Washington D.C.: U.S. Government Printing Office).
- van Praag, B., A. Hagenars and J. van Weeren (1982) 'Poverty in Europe,' *Review of Income and Wealth*. Vol 28.
- World Bank (1994) *Uzbekistan: Adjusting Social Protection*, Report No. 13023-UZ.
- World Bank (1995a) *The Kyrgyz Republic: Poverty Assessment and Strategy*, Report No. 14380-KG.
- World Bank (1995b) *Pakistan Poverty Assessment*, Report No. 14397-PAK.