WHAT IS POVERTY AND WHO ARE THE POOR? REDEFINITION FOR THE UNITED STATES IN THE 1990'S[†]

Absolute versus Relative Poverty

By JAMES E. FOSTER*

Should poverty be measured using an "absolute" or a "relative" approach? This age-old question in poverty measurement is once again on the agenda, due to the ambitious proposals of Patricia Ruggles (1990) and the National Research Council of the National Academy of Sciences (Constance Citro and Robert Michael, 1995) to alter the way U.S. poverty is measured. Their wide-ranging suggestions include a new "hybrid" approach to setting the poverty threshold that, unlike the current absolute method, is sensitive to changes in the general living standard, but less sensitive than a purely relative approach. The proposals also recommend using aggregate indexes of poverty beyond the usual "head-counts," such as well-known "gap" measures and indicators of the distribution of resources among the poor. Important relative notions of poverty enter at this "aggregation" step as well. The effects of the various recommendations on the trend and cross-sectional profiles of poverty are actively being explored (see e.g., David Betson and Jennifer Warlick, 1997; Thesia Garner et al., 1997; David Johnson et al., 1997). At the same time it may prove useful to consider some of the conceptual measurement issues arising from the proposals. This is the direction taken in the present study.

This paper evaluates the multiple notions of relative and absolute poverty that arise in choosing poverty lines and in aggregating the data into an overall index of poverty. A general taxonomy is presented, and the question of robust comparisons is addressed within this general framework. Special attention is paid to distinguishing between (i) the general concept underlying the poverty line and (ii) the particular cutoff chosen. The paper concludes with a discussion of "hybrid" poverty lines and the associated parameter that is likely to play a key role in future discussions: the income elasticity of the poverty line.

I. Elements

Poverty measurement is based on a comparison of resources to needs. A person or family is identified as poor if its resources fall short of the poverty threshold. The data on families are then aggregated to obtain an overall view of poverty.

There are many ways of defining resources, constructing thresholds, and aggregating the resulting data (see e.g., Ruggles, 1990; Martin Ravallion, 1994; Citro and Michael, 1995). Virtually all partition the population into groups of families (or resource-sharing units) with similar characteristics, and I follow this approach here. Let Θ denote the raw data, containing information on resources received by families, their demographic and other characteristics, and perhaps other data (e.g., consumption distributions) needed to construct poverty thresholds. Let m be the number of distinct groups, with $n^k = n^k(\Theta)$ being the number of families in group k. Once a specific definition of family resources has been fixed, this yields a distribution of resources among the families in group k, denoted by the n^k-dimensional vector $\mathbf{x}^k = \mathbf{x}^k(\Theta)$. The poverty threshold for families in group k is denoted by the number $z^k = z^k(\Theta)$; a family is identified

[†] *Discussants:* David S. Johnson, U.S. Bureau of Labor Statistics; Patricia Ruggles, U.S. Department of Health and Human Services; Barbara Wolfe, University of Wisconsin; Christopher Jencks, Harvard University.

* Department of Economics, Vanderbilt University, Nashville, TN 37235. I thank the discussant, David S. Johnson, for his insightful commentary. Financial support from the John D. and Catherine T. MacArthur Foundation through the Network on Inequality and Poverty in Broader Perspective is gratefully acknowledged.

as poor if its resource level falls below z^k. Exactly how z^k is to be set (i.e., the "identification step" of Amartya Sen, 1976) is a key part of the present discussion.

As for the "aggregation step," most U.S. studies report poverty levels for the demographic groups and then aggregate to obtain an overall level of poverty. Thus, they implicitly take the poverty index to be "decomposable" across the groups (on which more will be said presently). With overall poverty a weighted sum of group poverties, the aggregation question reduces to a choice of the poverty index $P(\mathbf{x}; z)$ to apply to a typical group distribution $\mathbf{x}$ and poverty line z. The most common index is the head-count ratio $H(\mathbf{x}; z) = q/n$ where q is the number of poor families in $\mathbf{x}$ given z, and n is the number of families in $\mathbf{x}$. This index provides important information on poverty (namely, the frequency of poverty among the population) but ignores other relevant information on the depth and distribution of poverty. Another important kind of "partial index" is based on the sum of the income gaps $(z - x_i)$ of poor families. These "gap indexes" add a second dimension of "depth" to poverty evaluations. A third dimension is provided by indexes of inequality among the poor.

While each partial index conveys useful information about some aspect of poverty (assuming, of course, that the poverty threshold itself is meaningful), one must be careful in using its unidimensional prescriptions as a guide to policy. For this and other reasons, it has been argued (see e.g., Foster and Sen, 1997) that an index combining all three dimensions is more coherent in this role. Such "distribution-sensitive" indexes have been used to great advantage in international comparisons and development studies (see e.g., Ravallion, 1994).

II. Absolutes and Relativities

There are several ways in which relative and absolute considerations enter into poverty measurement. I offer a simple taxonomy including the *threshold* and *equivalence-scale* choices in the identification step, and the treatment of *population, scale,* and *individual deprivation* in the aggregation step.

A. *Threshold*

The first and perhaps most important sense in which poverty measurement is absolute or relative concerns the setting of the poverty standard. An *absolute* poverty line is a fixed (group-specific) cutoff level z_a that is applied across all potential resource distributions. In comparisons over time, for example, the standard is unchanged even in the face of economic growth (although provisions are made for changes in price levels);[1] similarly, in comparisons across countries, fixed-threshold comparisons require an appropriate exchange rate. If the absolute standard is truly independent of the current data, though, how can one be sure that the standard chosen is an appropriate one? The poverty line is typically calibrated in some initial period using, say, food-budget studies, and it is then carried forth from year to year, irrespective of whether the same procedure applied to current data would yield the same result. In a growing economy, the gap between the hypothetical recalibrated level and the historical standard may well be quite large. Such is the case with the current U.S. poverty standard, and this is one of the criticisms that have been leveled against it (see Citro and Michael, 1995 pp. 2–3).

In contrast, a relative approach uses current data to generate the current poverty threshold. A *relative* poverty line begins with some notion of a *standard of living* $r(\mathbf{x})$ for the distribution $\mathbf{x}$, such as the mean, median, or some other quantile, and defines the cutoff as some percentage α of this standard. The result is a poverty threshold $z_r = \alpha r(\mathbf{x})$ that varies one-for-one with the standard of living, in that a 1-percent increase in r is matched by a 1-percent increase in z_r. Examples include the "50 percent of the median" relative poverty line proposed by Victor Fuchs (1969) and the "50 percent of the mean" threshold employed by

[1] There is a significant issue of whether resource should be expressed in real terms and, if so, which cost of living index to use. This issue is ignored here for simplicity, but it is clearly another potentially important source of "relativity" in the measurement of poverty.

Michael O'Higgins and Stephen Jenkins (1990).[2]

Using a relative line does not amount to measuring inequality (although theorem 6 in Foster and Anthony F. Shorrocks [1988a] provides one important link) nor does it imply that poverty is by definition "always with us" (see Anthony Atkinson, 1975 p. 189). And while many studies regard absolute lines as being especially low and relative lines as being high, this is not necessarily the case. If living standards are rising and thresholds are pegged at $z_a = z_r$ in some initial period, then $z_a < z_r$ for all subsequent periods, but $z_a > z_r$ for all previous periods, as emphasized by Citro and Michael (1995 p. 132). In any isolated period, it is not possible to tell whether a given threshold z is relative or absolute, nor is the distinction particularly important, since the same numerical cutoff, however originally derived, must lead to the same level of poverty.

The key distinction between absolute and relative thresholds is not seen in the specific values obtained at a given date, but in how the values change as the distribution changes. Thus, there is an important distinction to be made between the general *concept* underlying the poverty threshold, and the specific *cutoff* selected. For comparisons involving extended periods of time, or very different standards of living, the former is likely to be the more important issue (see also Ruggles, 1990 Ch. 3), while the latter choice (of cutoff) is largely arbitrary (see Fuchs, 1969; Atkinson, 1975, 1987; Foster and Shorrocks, 1988b). This inevitable arbitrariness casts doubt on the meaning of the cardinal poverty levels obtained at specific cutoffs and leads to a consideration of the robustness of results to changes in the cutoff, a topic I will return to below.

B. *Equivalence Scale*

A second entry point for relativities in poverty measurement is where poverty lines are adjusted across demographic groups. One approach is to apply repeatedly the procedure for setting poverty lines to each group separately and thereby arrive at m independent thresholds. However, as noted by Ruggles (1990 Ch. 4), this can lead to odd (nonmonotonic) behavior of the poverty line as family size changes. An alternative approach sets the line in one reference group and then derives the remaining thresholds using an "equivalence scale" to account for the differing needs of different-sized families. The typical scale provides the rate at which a dollar for one group translates into dollars for another. So if group 1 is the reference, and s^k is the conversion rate from group 1 to group k, then $z^k = s^k z^1$ becomes the poverty line for group k.

This sort of equivalence scale is *relative* in that the transformation from group to group is multiplicative, and consequently group poverty lines are proportionate to each other. Another possibility raised by Charles Blackorby and David Donaldson (1994) is for variations in family configuration to have an constant *absolute* effect so that, for example, adding another child is seen as an additional fixed (real) cost to the family, independent of the size of the base threshold. Relative equivalence scales preserve the ratios of group poverty lines as the base threshold changes; an absolute equivalence scale preserves the absolute differences. The two forms are indistinguishable for a single observation or if the reference threshold remains unchanged (as with an absolute poverty line).

C. *Population*

The aggregation stage uses three notions of absolute and relative poverty in constructing poverty indexes. First, a *relative* or per capita poverty index is independent of the population size in the sense that "replicating" the population leaves the poverty value unaffected: for example, $P(\mathbf{x}, \mathbf{x}; z) = P(\mathbf{x}; z)$. Such a measure is based purely on the relative frequencies of incomes in the income distribution. In contrast, an *absolute* index is one whose value rises in proportion to the number of replications: for example, $P(\mathbf{x}, \mathbf{x}; z) = 2P(\mathbf{x}; z)$. The head-count ratio q/n is relative in this sense while the head-count q is absolute. An

[2] There are important measurement issues in selecting the standard of living. Should it be the mean, the median, or some other representative income? Should it be from the entire population or some reference group? Should it be for all expenditures or a significant subset? (For references, see Citro and Michael [1995].)

absolute index can be converted to a relative index by dividing by n.

D. *Scale*

A second notion concerns the behavior of an index when the poverty line and incomes are simultaneously altered. A *relative* or scale-invariant index is one that is unchanged when the poverty line and all incomes are multiplied by the same factor. An *absolute* or translation-invariant index is independent of additions of the same constant to the poverty line and all incomes (Blackorby and Donaldson, 1980). Thus, for example, the aggregate poverty gap $\Sigma_{i=1}^{q}(z - x_i)$ is an absolute index of this sort, while the normalized poverty gap $\Sigma_{i=1}^{q}(z - x_i)/z$ (which measures the poverty gap in poverty line units) is relative. The head-count ratio is both absolute and relative in this sense (and is essentially unique in this respect [see Buhong Zheng, 1994]).

E. *Deprivation*

Finally, the basic notion of deprivation that underlies a given index may be relative or absolute. If a family's poverty level depends purely upon its own characteristics, its resource level, and its threshold, then the index is based on a notion of *absolute deprivation.* Foster and Shorrocks (1991) relate this to *decomposability* of the index across population subgroups (overall poverty is a weighted sum of subgroup poverties for any partition) and also to a more fundamental notion of *subgroup consistency* (overall poverty is increasing in subgroup poverty levels for any partition). The head-count ratio and the gap indexes are absolute in this sense, as is the index of Foster et al. (1984) which takes $[(z - x_i)/z]^2$ as the ith poor family's deprivation level. In contrast, Sen's (1976) index is founded on the notion of *relative deprivation,* since a family's deprivation level depends crucially on its relative position among poor families and thus incorporates information beyond its own data. A discussion of the two approaches can be found in section A6 of Foster and Sen (1997).

III. Robust Comparisons

The above taxonomy presents several avenues for relative and absolute concepts to enter into poverty evaluations, and many combinations are possible. For example, the current method for evaluating U.S. poverty employs an absolute threshold for each group and a relative *or* absolute equivalence scale (indeterminate since poverty lines are unchanging) to identify the poor; for the aggregation step it typically uses the head-count ratio, a population relative index that is both absolute and relative with respect to scale, and which is based purely on absolute deprivation. The aggregate gap, which is absolute in all three dimensions, is often used as an alternative index.

Each combination of absolutes and relative concepts has many possible implementations (i.e., specific cutoffs, scales, and indexes) from which to choose. Inevitably, this entails making choices for which there is little guidance (why 50 percent of the median instead of 49 percent?). It is important to note, however, that the decision need not be based on normative or subjective considerations. The selection from the array of possible implementations could be purely arbitrary—made in the interest of getting on with the analysis (on this distinction, see Sen [1980]).

Given the inherent arbitrariness in selecting a specification, it is important to evaluate the robustness of any conclusions obtained. In cases where the numerical poverty levels are important, this may be as simple as testing other reasonable specifications and reporting how the poverty level changes. Betson and Warlick (1997), for example, use 20-percent changes in z to illustrate the cardinal sensitivity of head-counts to the threshold. Alternatively, when rankings of poverty levels are all that matter, one has available a rather large collection of tools to evaluate ordinal robustness (analogous to the well-known Lorenz criterion for inequality analysis), which cover variable thresholds, equivalence scales, and indexes (see e.g., Foster, 1984; Foster and Shorrocks, 1988b; Atkinson, 1987, 1992). Virtually all approaches trace back to notions of stochastic dominance from risk theory (see the general discussion in Foster and Sen [1997]).

Most results of this type are presented in a one-group framework with absolute thresholds; but in fact, the tools have far greater applicability. As an illustration, suppose that the base threshold z^1 and equivalence scale s^k are relative, the index P is based on a notion of absolute deprivation (hence decomposable) but otherwise relative, and the only question is the specific cutoff α to be used in setting the relative poverty line. Suppose that for a specific value of α, say, $\alpha = 50$ percent, the resource distribution $(\mathbf{x}^1, \ldots, \mathbf{x}^m)$ has greater poverty than $(\mathbf{y}^1, \ldots, \mathbf{y}^m)$. When can one be sure that this will remain true for an entire range of α values, say $(0, \bar{\alpha})$ where $\bar{\alpha} > 50$ percent? Let r be the standard of living underlying the relative poverty line, and let r_x and r_y denote the respective standards in the distributions. Construct a new "equivalent" distribution $\tilde{\mathbf{x}}^k$ for demographic group k by dividing family resources by the equivalence scale s^k, and then replicating by family size in k, so that $\tilde{\mathbf{x}}^k$ has one equivalent resource level for each person in group k. It is not difficult to show that for P satisfying the above properties, the poverty level of the original distribution $(\mathbf{x}^1, \ldots, \mathbf{x}^m)$ at the group-specific thresholds $z^k = s^k z^1$ is simply

$$P(\tilde{\mathbf{x}}^1, \ldots, \tilde{\mathbf{x}}^m; z^1)$$

or the poverty in the equivalent distribution given group 1's poverty line. If one further normalizes incomes by the standard of living, then the poverty level is given by

$$P(\tilde{\mathbf{x}}^1/r_x, \ldots, \tilde{\mathbf{x}}^m/r_x; \alpha).$$

Consequently, the judgment that $(\mathbf{x}^1, \ldots, \mathbf{x}^m)$ has greater poverty than $(\mathbf{y}^1, \ldots, \mathbf{y}^m)$ is in fact robust in α if

$$P(\tilde{\mathbf{x}}^1/r_x, \ldots, \tilde{\mathbf{x}}^m/r_x; \alpha)$$

$$> P(\tilde{\mathbf{y}}^1/r_y, \ldots, \tilde{\mathbf{y}}^m/r_y; \alpha)$$

for all $\alpha \in (0, \bar{\alpha})$.

This last condition is in a form that allows the application of results in Foster and Shorrocks (1988b) and Atkinson (1987). So, for example, the test for the head-count ratio H checks whether the two distributions of normalized equivalent incomes can be compared using first-degree stochastic dominance over the range $(0, \bar{\alpha})$, while the tests for the normalized gap index and the Foster et al. (1984) index use second- and third-degree stochastic dominance, respectively. Atkinson's (1987) results go beyond these results to consider variations in poverty indexes and indicate, for example, that if there is an unambiguous comparison for H (and hence first-degree stochastic dominance), then virtually any acceptable index P will agree with this conclusion. This illustrates the power of the head-count ratio in this context.

IV. Hybrid Measurement

Many of the categories in my taxonomy allow for an intermediate position to be chosen in place of a pure relative or absolute approach. One particularly interesting example is the "hybrid" poverty threshold that is central to the proposal in Citro and Michael (1995), which is based on what might be termed a "partial" standard of living: r_p is the median expenditure on certain basic goods. The threshold $z = \alpha r_p$ has the same structure as a purely relative cutoff (and in fact the robustness result applies equally well to it). However, median expenditures on basic goods do not rise as fast as, say, median total expenditures, and it is this empirical fact that gives z its hybrid nature.

One could also imagine thresholds that are hybrid by construction, in that they depend directly on an absolute and a relative standard. For example, consider a weighted geometric average of a relative threshold $z_r = \alpha r$ and an absolute threshold z_a, namely, $z = z_r^{\rho} z_a^{1-\rho}$, where $0 < \rho < 1$. This form of hybrid line has the property that a 1-percent increase in the living standard r always leads to a ρ-percent increase in the poverty line. In other words, ρ is the elasticity of the poverty line with respect to the living standard, or what Gordon Fisher (1995) has termed the *income elasticity of the poverty line*. In general, $\rho = (dz/dr)(r/z)$ has a natural interpretation as a measure of the extent to which a given threshold z is relative, with $\rho = 0$ corresponding to an absolute poverty line and $\rho = 1$ a fully relative one. The possibility of using a hybrid standard changes

the question "absolute or relative?" to "exactly how relative?" with ρ as the relevant decision variable.

In his defense of the relative approach, Fuchs (1969 p. 201) posited that the cutoff "would be recognized as a national value judgment and would be arrived at through the normal political process." One theme of the present paper is the primacy of general concept over specific cutoff; if this is accepted, then the subject of public discourse would more properly be ρ, the income elasticity of the poverty line. The choice of ρ would then answer the normative question: "To what extent should the poor share in economic growth?" An elasticity of 1 appears to be too high to command much political support in the United States. An elasticity of 0 is implicit in the current standard, but given the historical tendency for absolute standards to be periodically revised (Fisher, 1995) and the long-standing explanations why, when the general standard of living rises, resources may need to be higher to achieve the same ends (e.g., Atkinson, 1975; Sen, 1983), this answer may not be tenable in the long run. However, it remains to be seen whether the particular hybrid standard proposed by National Research Council of the National Academy of Sciences, which has a historical income elasticity of $\rho = 0.65$ (Citro and Michael, 1995 p. 143), will garner enough support to displace the current standard.

REFERENCES

Atkinson, Anthony B. *The economics of inequality*. Oxford: Oxford University Press, 1975.

_____. "On the Measurement of Poverty." *Econometrica*, July 1987, *55*(4), pp. 749–64.

_____. "Measuring Poverty and Differences in Family Composition." *Economica*, February 1992, *59*(233), pp. 1–16.

Betson, David M. and Warlick, Jennifer L. "Alternative Historical Trends in Poverty." Mimeo, University of Notre Dame, 1997.

Blackorby, Charles and Donaldson, David. "Ethical Indices for the Measurement of Poverty." *Econometrica*, May 1980, *48*(4), pp. 1053–60.

_____. "Measuring the Cost of Children: A Theoretical Framework," in Richard Blundell, Ian Preston, and Ian Walker, eds., *The measurement of household welfare*. Cambridge: Cambridge University Press, 1994, pp. 51–69.

Citro, Constance F. and Michael, Robert T. *Measuring poverty: A new approach*. Washington, DC: National Academy Press, 1995.

Fisher, Gordon M. "Is There Such a Thing as an Absolute Poverty Line Over Time?" Mimeo, U.S. Department of Health and Human Services, Washington, DC, August 1995.

Foster, James E. "On Economic Poverty," in Robert Basmann and George Rhodes, eds., *Advances in econometrics*, Vol. 3. Greenwich, CT: JAI Press, 1984, pp. 215–51.

Foster, James E.; Greer, Joel and Thorbecke, Erik. "A Class of Decomposable Poverty Measures." *Econometrica*, May 1984, *52*(3), pp. 761–66.

Foster, James E. and Sen, Amartya. "*On Economic Inequality* after a Quarter Century," in Amartya Sen, ed., *On economic inequality*. Oxford: Clarendon, 1997, pp. 107–219.

Foster, James E. and Shorrocks, Anthony F. "Inequality and Poverty Orderings." *European Economic Review*, March 1988a, *32*(2–3), pp. 654–61.

_____. "Poverty Orderings and Welfare Dominance." *Social Choice and Welfare*, 1988b, *5*(2/3), pp. 179–98.

_____. "Subgroup Consistent Poverty Measures." *Econometrica*, May 1991, *59*(3), pp. 687–709.

Fuchs, Victor. "Comment on Measuring the Size of the Low-Income Population," in Lee Soltow, ed., *Six papers on the size distribution of wealth and income*. New York: National Bureau of Economic Research, 1969, pp. 198–202.

Garner, Thesia I.; Paulin, Geoffrey; Shipp, Stephanie; Short, Kathleen and Nelson, Chuck. "Experimental Poverty Measurement for the 1990's." Mimeo, Bureau of Labor Statistics, Washington, DC, 1997.

Johnson, David; Shipp, Stephanie and Garner, Thesia. "Developing Poverty Thresholds Using Expenditure Data." Unpublished manuscript presented at the Joint Statistical Meetings, Anaheim, CA, August 1997.

O'Higgins, Michael and Jenkins, Stephen. "Poverty in the EC: Estimates for 1975, 1980, and 1985," in Rudolph Teekens and Bernard M. S. van Praag, eds., *Analysing poverty in the European Community: Policy issues, research options, and data sources*. Luxembourg: Office of Official Publications of the European Communities, 1990, pp. 187–212.

Ravallion, Martin. *Poverty comparisons*. Chur, Switzerland: Harwood, 1994.

Ruggles, Patricia. *Drawing the line*. Washington, DC: Urban Institute Press, 1990.

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

_____. "Description as Choice." *Oxford Economic Papers*, November 1980, *32*(3), pp. 353–69.

_____. "Poor, Relatively Speaking." *Oxford Economic Papers*, July 1983, *35*(2), pp. 153–69.

Zheng, Buhong. "Can a Poverty Index Be Both Relative and Absolute?" *Econometrica*, November 1994, *62*(6), pp. 1453–58.