Counting the world's poor: problems and possible solutions

Angus Deaton
Research Program in Development Studies
Princeton University

August, 2000
Revised, December 2000

0. Introduction

The World Bank prepares and publishes estimates of the number of poor people in the world.

While everyone knows that these numbers should be taken with a pinch of salt, the numbers are

arguably important. In an institution where the reduction of poverty is the paramount objective,

some overall yardstick of progress (or the lack of it) is required. The numbers are frequently

quoted by politicians and by leaders of international organizations, including the Bank itself, who

believe the numbers are effective for advocacy. Indeed, there is a long history of studies of

poverty mobilizing support among the non-poor for anti-poverty policies. So it is important to

know whether the world and national poverty counts are sound enough to support these uses. As

recent discussions have made clear, the apparent lack of poverty reduction in the face of

historically high rates of economic growth, both in the world as a whole, and in specific countries

(most notably India), is providing fuel for the argument that economic growth does little to

reduce poverty. How confident can we be that the data actually support these inferences? Are the

changes in the poverty counts sufficiently well-measured to support conclusions about growth

and poverty reduction? Should the World Bank stand ready to be judged by its success in

reducing the current measures of world or even national poverty? If not, can better data collection

or better methodologies improve the numbers?

A more fundamental issue is whether the number of people suffering consumption (or

income) deprivation is the right yardstick. It is widely agreed that other dimensions of poverty are

as, or more, important than income poverty, particularly deprivations in health, education, and

democratic rights. There are difficult measurement issues associated with all of these other

dimensions, especially if we try to combine them into a single measure, such as the Human

1

Development Index. Recall too that, for many purposes, we need to monitor changes over periods of a few years, which is typically quite difficult for non-income based measures. For example, the average literacy rate of a population changes too slowly to be much use as a monitoring device in all but the longest of runs. Even so, and independently of measurement issues, the importance of these non-income dimensions of poverty should be a warning not to overemphasize income poverty at the expense of other measures. It should also be noted that other dimensions of deprivation are just that, "other" measures. They are not perfectly correlated with income or consumption poverty so that, just as we cannot subsume health or literacy measures into an income measure, we also should not try to measure income poverty by using health or education as a proxy. Difficulty in measuring income poverty is a poor argument for focusing on other dimensions of welfare and deprivation.

This paper discusses a number of problems with the current $1-a-day poverty counts, makes some suggestions for improvement, and identifies issues that need further research. World poverty numbers are calculated in two stages. At the first, or international stage, a world poverty line is set and used to derive comparable poverty lines for each country. At the second, or domestic stage, the poverty lines are used to count the number of poor people in each country, and the totals added up over countries. It is useful to think about each of these two stages separately, and there is disquieting evidence about both. At the international level, the purchasing power parity exchange rates that are used to turn the $1-a-day into national currencies are arguably inappropriate in theory, and in practice, their regular revision (to different base years with different relative prices) plays havoc with the poverty estimates, changing them in ways that have little or nothing to do with the actual experience of the poor. Within each country, at the

domestic stage, the Bank's poverty count comes from household surveys, and the incorporation of these unit record data into its statistical base is surely one of the Bank's major achievements over the last twenty years. Yet, in many countries, including India where around a third of the world's poor live, there are large and growing discrepancies between the survey data and the national accounts. Because poverty counts come from the survey data, while growth measures come from the national accounts, and because they are evidently measuring different things, there is no consistent empirical basis for conclusions about the extent to which growth reduces poverty. That economic growth, *as measured,* has at best a weak relationship with poverty, *as measured*, means little more than would a finding that growth in China had failed to reduce poverty in India.

I discuss the international issues first, followed by the domestic ones. For those who are unpersuaded that the world poverty counts are useful, or believe them to be "essentially propagandist" (Srinivasan, 2000), the international issues are irrelevant. But the domestic problems are important not only for the world count, but for anyone who cares about tracking income poverty within individual countries. A final section summarizes and concludes, and lays out the main research issues that require attention.

1. International

1.1 Purchasing power parities and the base poverty line

The current concept for world poverty is the number of people who live in households whose daily consumption per head is less than the purchasing power parity (PPP) equivalent of $1 a day in constant 1985 PPP dollars. In the latest revision, this has been updated to $1.08 in constant

1993 PPP dollars, but is still conveniently referred to as the $1 a day poverty line, see Chen and Ravallion (henceforth CR) (2000). The virtues of this approach are many. It is simple, easy to remember, and applies equally to all countries. It is denominated in a currency that is familiar to the relatively wealthy people who are the primary users of the measures, and who are the primary target for rhetoric based on them. The $1-a-day was originally selected as being representative of poverty lines in use in low-income countries, and the new $1.08 is representative of current lines, and is thus anchored in actual practice. The $1.08 a day is converted to the 1993 value of local currencies using internal Bank PPP tables, and then updated to the target year using a local consumption deflator. The resulting poverty line is then used to count the number of people in each poor (technically World Bank Part 2) country using household survey data, when possible from the raw data and, when not, from interpolations using published tables from such surveys.

Given world—and U.S.—inflation between 1985 and 1993, it is somewhat surprising that the international poverty line should have increased by only 8 percent, from $1 to $1.08. But the updating was carried out by going back to the country poverty lines, and converting back to international dollars, so that the modest increase comes, not from a failure to allow for world inflation, but because the PPP international dollar has strengthened relative to the currencies of the poor countries whose poverty lines are incorporated into the international line.

The PPP numbers come from data collected from 110 countries in 1993 by the International Comparison Project. Earlier PPP numbers came from the Penn World Tables, which covered only 60 countries. The revision from the earlier PPPs to the new PPPs in the most recent revision of the poverty counts, CR (2000), resulted in some disconcertingly large changes in poverty counts even for the same country in the same year, and using the same survey data. The worst

4

cases are Sub-Saharan Africa, where the poverty rate in 1993 (i.e. the *same year*) rose to 49.7 percent from 39.1 percent, Latin America, with poverty rates falling to 15.3 percent from 23.5 percent, and Middle-East and North Africa where the poverty rate fell from 4.1 percent to 1.9 percent. Changes of this size risk swamping real changes, and it seems impossible to make statements about changes in world poverty when the ground underneath one's feet is changing in this sort of way. Note that CR (2000) recalculate previous counts at the new PPPs, so the issue is *not* that we are comparing current poverty using one PPP conversion with earlier poverty using another conversion. The problem is that, if revisions can have such enormous effects, it is difficult to place much confidence in any set of poverty estimates, either of levels or of changes, even those based on the latest and best set of PPP exchange rates.

For countries whose poverty lines are near the middle of the distribution of consumption or income, poverty rates can be extremely sensitive to changes in the poverty line, including those induced by PPP (or other) errors. The 1997 *World Development Report* listed Thailand as having only 0.1 percent of its population as less than $1 a day at PPP; this number was cited in the New York Times by then Chief Economist Joseph Stiglitz as one of the consequences of the Asian Economic Miracle, but is much more likely a tribute to inappropriate PPP conversion. This anecdote suggests that, at a minimum, it would be wise to check the PPP poverty lines for plausibility *in the countries* to which they apply. Of course, if the lines are wildly implausible, we need some other method, a issue to which I return below.

Revisions to PPP exchange rates are of two, conceptually different kinds. First, there are the usual revisions that come with better information, elimination of previous errors, and so on. For example, it is claimed that for China, the new updated PPP rates are much more solidly based

than were the old. There are also 50 more countries for whom PPP exchange rates are calculated, rather than imputed. Second, there are the revisions that come from the change of base given that PPP exchange rates depend on the relative prices of commodities in the base year. This second source of revision ultimately traces back to the conceptual, or "index number" problems associated with making international comparisons of well-being. Indeed, one might argue that the $1-a-day poverty line, by its very simplicity, is misleading just because it seems to sweep away the problems. The calculation of the cost of subsistence for a poor Indian versus that for a poor Kenyan or Brazilian is fraught with conceptual difficulties. In particular, if we take the Indian as reference, and calculate her relative costs of living in Calcutta versus (the hard to imagine) costs of living after being moved to Sao Paulo, then we are likely to get a very different number from doing the same with the Brazilian transplanted to Calcutta. Suppose that, in country A, the staple food for poor people is beans, and in country B, it is rice. If we price beans in country B in B's currency, and divide by the bean price in A in A's currency, we get an implicit "beans" PPP exchange rate. A similar calculation for rice gives a "rice" PPP exchange rate. These will give the same answer (or at least reciprocals of one another) only if both goods are freely traded between the two countries without tariffs, barriers, or transportation costs, conditions that are far from holding true in practice, even approximately. Many (perhaps most) of the world's poor eat food staples that were grown only a short distance away, and whose price is only loosely (if at all) connected to world commodity prices. And the staple food in one country (or even region) is often a relatively rare (and even expensive) luxury somewhere else (e.g. sticky rice in Northern Thailand versus the United States), so that the exchange rate obtained from the prices of the same commodity is not very useful for converting poverty lines from one country to another.

6

In practice, matters are not quite so stark, and actual PPPs price a representative bundle of goods in each country, and compare the local cost of the bundle with the US dollar cost of the same bundle. But note that the answer depends on the structure of *relative* prices at the time of the comparison so that, when PPPs are revised with a new base year, the new exchange rates are not simply a new measure of the *old* concept, but a new measure of a *new* concept. World prices of primary commodities are notoriously volatile, and for some countries, primary commodities are a large share of their GDP. As a result, the PPP exchange rate for Nigeria (for example) relative to the US will vary with the world price of oil in the base year. Which is presumably why the African and Latin American poverty rates are so sensitive to revisions of the PPP base year. And while one might argue that the changes in the world prices of beans and rice ought to change the relative poverty lines of Brazilians and Indians, it is much harder to make the case for changes in the world price of oil. Even in theory, PPP exchange rates as currently defined are not designed to convert poverty bundles, see also Srinivasan (2000).

One way of improving the PPP poverty lines would be as follows. First best would be a International Comparison Program that focuses, not on GDP and its components, but on some more appropriate poverty bundle. Though in this case the first-best is perhaps not very good, if only because even a poverty-bundle PPP would not address the underlying index number issues. A second-best short-cut would be to take the current set of 1993 poverty lines, updated to the present, and to check them in each country (or at least those countries that have significant numbers of poor people.) In cases where the numbers are implausible (for example, showing no poverty in Thailand,) the PPP exchange rate can be calculated using the first-best procedure, by pricing out in dollars and in local currency an appropriate poverty bundle, for example the

7

average consumption bundle of the bottom quintile of the population. While the choice of the

bottom quintile is arbitrary (although sensible), the calculations are straightforward given a

household expenditure survey, which often collects data on quantities as well as on expenditures.

When quantity data are not in the survey, they can be calculated from the expenditure data using

a set of local consumer price quotes, such as those that are routinely used for the construction of

national consumer price indexes.

Armed with a set of PPP poverty lines that have been locally validated, and adjusted as

necessary, the crucial next step is to hold these fixed (in real terms), and not make further

adjustments as PPP exchange rates are revised in response to changes in the base year and its

associated relative prices. The benefit of this proposal is that African and Latin American poverty

rates are not held hostage to fluctuations in the prices of primary commodities in the base year of

PPP calculation, and that world poverty rates are consistently calculated against a fixed target. Of

course, it could be argued that the line is no longer the constant PPP $1-a-day line, and that the

Bank is failing to update for new information. But this objection is more semantic than real. The

lines originally started from $1-a-day, and can retain the label, and they are updated by local

prices, just not for changes in world relative prices which effectively change the *definition* of PPP

exchange rates, and whose shifting basis is an embarrassment, not an advantage.

Even if the lines, once set, are to some extent arbitrary, and even if it were the case that a

reworking from scratch would lead to new lines, there is much to be said for holding them fixed.

Indeed it is hard not to be impressed by the durability of official poverty lines in countries as

diverse as India and the United States, whose respective poverty lines have survived unchanged

(except for inflation correction) for 30 and 40 years, respectively. Once an initial poverty line has

been set, and provided it commands public and political support at the time it was set, it seems that it can generate useful estimates of poverty for many years. Moreover, and for both India and the U.S., the original nutritional basis for the lines is often cited in their support. The same could be true for the $1-a-day lines.

1.2 International comparability without PPP

There are a number of alternatives to the $1-a-day PPP poverty line, though none is without is own problems. There is a tradition, followed in both the United States and India, of setting poverty lines with reference to what Ravallion (1998) refers to as "the nutritional requirements for good health." This is defined as the level of income (or total consumption) at which, on average, nutritional norms are met. Note that, because the standard is set in terms of income or total consumption, this is *not* the amount necessary to purchase the nutritional norms only, but the amount spent on the nutritional norms *and on other goods* by an average household who just meets the norms. An internationally comparable set of poverty lines could be set in this way with reference to a universal calorie norm, say 2,000 calories per person per day, and by using a national household expenditure survey for each country to calculate the level of total household expenditure per capita at which these norms are met in expectation.

Again, this method is not quite what it seems, and there are important practical qualifications. First, as with the recommendations for the amended PPP method, the procedure is used once and once only. The resulting poverty lines for each country are subsequently held fixed over time. This may seem contradictory; if the calorie method is correct, it should surely be applied consistently over time and space? But it is not correct, in the sense that it is a "scientifically

9

sound" method of setting the line. Instead, it is a device for generating a socially acceptable poverty line which can then be held fixed over time. The food rhetoric helps build legitimacy for the line, but judging from India and the United States, maintaining that legitimacy does not require updating, except possibly in the very long term. Successive commissions and Expert Groups have altered details of the Indian lines, for example, by introducing variation across states, see Government of India (1993), and in the US there have been frequent (sometimes detailed) discussions of reform, particularly National Research Council (1995), but neither the NRC's recommendations nor others have been adopted.

Second, it is important *not* to set separate nutritional poverty lines for different regions or sectors within a country. At the same level of per capita total household expenditure, urban people spend less on food, buy more expensive calories, and consume fewer of them, so that applying the method to urban and rural areas using the same calorie target will lead to higher poverty lines in urban than rural areas. Across regions with different income levels, the operation of Engel's Law has much the same effect, so that poverty lines can move more or less in proportion to average incomes, generating what are effectively relative not absolute lines, see Ravallion and Bidani (1994) and Ravallion and Sen (1996). Although relative lines may make a good deal of sense in other contexts, they are not appropriate for world counts that are based on an explicitly absolute standard.

2. Domestic issues in counting the poor

2.1 National Accounts versus Household Survey Data

A major source of current controversy is the discrepancy between national accounts and survey

estimates of consumption, especially when, as currently in India, the discrepancy is increasing. India is the largest single contributor to the world poverty count, with more than 400 million poor people out of a world total of 1.2 billion. The ratio of National Sample Survey (NSS) consumption to National Accounts Statistics (NAS) consumption has fallen from around unity in the 1950s, to 89 percent in 1968–69, to around 75 percent in the late 1970s to a little more than a half in recent years, see Bardhan (1974), Mukherjee and Chaterjee (1974), Srinivasan, Radhakrishnan, and Vaidyanathan (1974) for early discussions, and Srinivasan (2000) for a more recent update. Although different authors give somewhat different ratios, there is general agreement that the discrepancy is widening and that the difference in per capita growth rates is currently about 2 percent a year. As a result, the rapid post reform growth in GDP and consumption per head measured in the NAS does not show up in the NSS surveys. The latter are now the exclusive basis for the official poverty counts, and they show little or no reduction in official poverty during the reform period.

The Indian example is only the most notable of similar or related phenomena elsewhere, including some industrialized countries such as the US. Triplett (1997) estimates that, even after correction for differences in concept, between 1973 and 1994, per capita consumption in the national accounts has grown 0.4 percent a year faster than per capita consumption in the Consumer Expenditure Survey (CEX) and, that between 1984 and 1994, the difference in growth rates has been a full one percent a year. (Note that the CEX data in the US are not used to derive the national poverty counts.) Ravallion (2000b) notes a similar drift in other countries and regions, particularly in China (the second biggest single contributor to world poverty) as well as in Latin America where survey consumption growth is only 70 percent of NAS consumption

11

growth. In Sub-Saharan Africa, consumption growth in the surveys is less well linked to the NAS, but is also consistently slower. And while there is no drift in the transition economies of Eastern Europe and Central Asia it is only because there appears to be no relationship whatever between consumption growth in the NAS and in the surveys!

If the transition economies are ignored, Ravallion cannot reject the hypothesis that the coefficient of the growth rate of NAS consumption on survey consumption is unity in a regression on a cross-section of countries. Note however that, as is appropriate for his purpose, Ravallion's units are countries, not people, and that most of the poor people in the world live in countries or regions (East Asia, India, Sub-Saharan Africa) where survey growth rates of consumption are substantially less than national accounts growth rates. And even when the coefficient is insignificantly different from one, the point estimate is 0.836, indicating substantial slippage, and is lower still in East Asia (0.628), Sub-Saharan Africa (0.645), and South Asia (0.742).

That the drift is so common around the world, and generally in the same direction, suggests a common cause, and there is clearly an important research project trying to track it down. For example, might it be that the positive skewness in consumption, together with the plutocratic basis of the NAS versus the democratic basis of the surveys, results in a tendency for average consumption in the surveys to lie beneath the population mean, and to do so by more when inequality is increasing? Some limited calculations show this to be unlikely for reasonable parameters. There are other plausible links between increasing inequality and increasing bias, for example if the rich are more likely to be missed in household surveys, something that is widely suspected in India, and has been well-documented in the U.S., Groves and Couper (1998).

Because growth around the world is almost universally assessed from the National Accounts Statistics, and because the poverty counts come entirely from the survey data, the discrepancies between the two drive a wedge between measured growth and measured poverty reduction. If we were to measure both growth and poverty from the household survey data, the discrepancy would vanish; there has been little or no measured growth, and little or no measured poverty reduction. Similarly, if we were to follow Bhalla (2000a, b), and treat the national accounts estimate of mean consumption as the truth, using it to "correct" the survey data, the discrepancy would also vanish but in the other direction; there would be lots of growth and lots of poverty reduction. In the meanwhile, claiming that growth has done little to reduce poverty is comparable to saying that, in spite of rapid growth in China, poverty in India remained the same. The NAS and the surveys are (evidently) measuring different things.

Which is correct? We don't know, though it seems safe to say that there are almost certainly errors in both. There is a longstanding prejudice by many economists (although not in India) against surveys and in favor of national accounts that is probably without basis. Some points are worth noting:

(a) The two concepts are different, and in particular there are items in National Accounts consumption, perhaps most importantly consumption of non-profits and the imputed rental of owner occupied dwellings, that are not included in the surveys. In India, it is thought that as much as a half of the current discrepancy between the two estimates can be accounted for by implicit rents; it is also plausible that these rents are growing more rapidly than other items of consumption, thus contributing to the differential rates of growth.

(b) Although expenditure estimates from household surveys are sometimes used in

13

constructing National Accounts consumption for some commodities, most items are derived as residuals, so that errors and omissions elsewhere in the accounts are automatically absorbed into consumption. One major problem is consumption by businesses, which has to be estimated and subtracted from the production totals to get the consumption of households. Consumption of coarse grains by animals presents similar problems and is often estimated using a set of multipliers that are difficult to estimate with any precision, and that are infrequently updated. It is quite unclear that the NAS estimates of consumption should be treated as the gold standard to which the survey estimates should ideally correspond. Indeed, Bardhan (1974) writes that "it would be highly improper to judge the goodness of NSS estimates by pointing to their divergence (or lack of it) from such residual estimates."

(c) There are many well known problems with data collection from household surveys. Best practice consumption measures (as in India or Indonesia) use very long lists of specific items. There is good evidence, Deaton and Grosh (2000), that attempts to shortcut this process lead to underestimation of consumption, yet such questionnaires are widely used by many of the surveys that are used in the calculations. Survey questionnaires tend to change slowly over time, and may not keep up with the increase in consumption of goods and services that may have been unknown a decade or two ago.

(d) Consumption surveys miss some households through refusals. Surveys in developing countries typically report very low refusal rates (for example, zero in the Indian NSS), but it is often unclear the extent to which willing respondents have been substituted for unwilling ones. There is some evidence that wealthy households are most likely to refuse to cooperate; for example, it is often impossible for enumerators to gain access to gated communities. If so, it

14

might be conjectured that real income growth, by increasing the fraction of such people, will lead to an increasing underestimation of consumption. Against this, it might be argued that the very poor, especially those without fixed abode, are also missed in household surveys.

(e) In most surveys, different kinds of households have different probabilities of being included, either by design or because of problems in the field, as when some types of household are more prone to refuse their cooperation. As a result, the survey results must be weighted in order to give an accurate representation of the population as a whole, and the calculation of suitable weights depends on the availability of accurate, up-to-date information about the population. Although statistical agencies have a number of tools for ensuring accuracy, it would not be surprising if the weights become progressively more inaccurate as the latest census becomes outdated. That said, the divergence between survey and NAS estimates in India shows no evidence of the saw-tooth pattern around censuses that would be predicted by such an account.

(f) It is tempting to try to use the behavior of other measures of poverty, such as literacy or infant mortality, to judge the relative plausibility of the NAS on one hand and the surveys on the other. There is undoubtedly good research to be done along these lines. Even so, definitive results should not be expected. Health and literacy measures should not be regarded as proxies for income, if only because they frequently change in ways that have little or nothing to do with income. In consequence, we do not have the reliable link between the income and non-income measures of poverty that would allow us to infer the former from the latter. Much the same can be said of attempts to look at indicators like the share of the budget devoted to food. While it is true that the food share drops with increases in income, it is also influenced by much else,

15

including relative prices, demographic changes, and the distribution of income, to mention only the most obvious. As a result, factors (such as an increase in inequality) that are possible suspects in driving the discrepancy between the national accounts and the surveys, are also likely to affect the relationship between food and total expenditure. In consequence, it is difficult to know how much weight to attach to calculations such as that in Bhalla (2000a) that impute the growth in total expenditure from data on food availability. Perhaps more promising—in practice, if not in theory—are the measures of living standards derived by Filmer and Pritchett (1998) from principal components analyses of the various asset ownership measures that are included in Demographic and Health Surveys.

It is of course possible for poverty to remain constant or to increase in the face of positive growth in average consumption if people in the lower tail of the distribution are left behind, and most or all of the growth accrues to people above the poverty line. That growth should disproportionately favor the better-off is entirely plausible in some settings, as in India where the sources of growth have been outside of the agricultural sector where most poor people obtain their living. But, differences of definition apart, the NAS and survey estimates of consumption cannot both be right, because their means are different, and are systematically becoming more different over time. So whatever the role of increasing inequality in moderating the poverty reducing effects of growth, one of the reasons why measured growth has not reduced measured poverty is the growing statistical discrepancy between the two sets of data on which each is measured.

What to do? There is a school of thought in India among the pro-reformers, as in Bhalla (2000a), which argues that the NSS consumption figures should be scaled up so as to match the

NAS, and the rescaled survey estimates used to calculate poverty. (Indeed, this is how the official poverty counts used to be done in India, and no very convincing reason was ever given for the change following the recommendations in Government of India, 1993.) Bhalla (2000b) argues the "means from NAS, distribution from surveys" formula should be adopted for the world poverty accounts, and calculates that growth has indeed reduced poverty. But there is a serious lack of evidence for at least some of the assumptions that would validate such a calculation. In particular, all the following must be true: (a) the NAS estimates are correct, (b) the survey estimates of the mean are incorrect and (c), in spite of (b), the consumption levels of each household in the surveys are correct up to a multiplicative factor. The last is a real stretch. Most of the plausible accounts of what might have gone wrong with the surveys, such as a progressively larger number of refusals, involve a systematic underestimation of inequality, as well as of the mean. It seems particularly implausible that the degree of underestimation in the surveys should be the same for the urban and rural households, see Ravallion (2000a). Indeed, many people in India, particularly those skeptical of the benefits of the reforms, believe that such growth as has taken place in consumption has accrued to people well-above the poverty line, and to urban more than rural people, and that those are the people who are being systematically under-sampled by the surveys. If so, one might suppose that the NAS modestly overestimates the rate of growth of consumption, that the NSS greatly underestimates it by missing the wealthy, but that the NSS-based poverty figures are essentially correct. And since the poor are typically not owners of valuable housing, the inclusion of an allowance for the rental value of owner occupied housing is unlikely to conclusion.

A good statistical principle is that, if two sources of data disagree, and we have no reason to

17

favor one over the other, then we should combine them to make a better estimate. This principle would argue for a more modest version of the scaling proposal, in which the survey data are scaled up by some weighted average of the NAS and the survey means, at least after correction for conceptual differences and coverage. Of course, the data on economic growth would also have to be scaled down accordingly. This proposal would make neither side happy, each of which has its own reasons for believing that one set of accounts is biased and the other bias free, but I can see no very convincing argument against it. Its limits are obvious; it takes no account of reasonable suspicions that the NSS errors are not neutral between rich and poor, or between urban and rural. But in the absence of the sort of better information that one might hope will be produced in the longer term, it is surely better than either of the extreme alternatives that are currently dominating the debate.

The urgent need here is to start a serious program of reconciliation between the NAS and survey data in a few countries, of which India must surely be one. (The Bureau of Labor Statistics and the Bureau of Economic Analysis at Census are talking about setting up a similar program in the US, and perhaps common cause could be made.) This program must work both with the National Accounts and with the survey organizations, and not start off by assuming one is right and the other wrong. It might be productive to start with a particular commodity, or group of commodities, such as cereals, where the production, export, and import data are relatively strong, and where there is a great deal of experience with household surveys. It would be possible to mount special supplementary household surveys for the task, which need focus on only a few goods. But it would also be a good idea to audit the performance of the survey agencies, just as the NAS agency is being audited. For example, the Indian NSS led the world in household survey

methodology and practice. But that was a long time ago, and the organization is no longer as dynamic and open to new practice as it once was. There has been a great deal of progress in survey methodology in the last 20 years, with the result that those statistical offices that used to be the leaders, mostly in South and East Asia, are now the most likely to be behind.

The Bank is the only international body that has the capacity and organization to run a research program designed to resolve the discrepancies between National Accounts and survey data. To its credit, the Bank has played a major role, through the LSMS surveys, in collecting household survey data on a nearly worldwide basis. The next step is to find out why these surveys give such different results from the National Accounts. Until it does so, neither the Bank nor anyone else will have a firm platform on which to make statements about the central issue of how growth affects poverty.

2.2 Price indexes

Local consumer price indexes are used to update a real poverty line for comparison with the nominal expenditures or incomes collected in the surveys. Consumer price indexes are of variable quality around the world and poverty counts can be very sensitive to any errors. Consumer price indexes are constructed from two components, a set of prices, collected on a regular schedule from retail shops and markets around the country, and a set of weights, which typically come from a household expenditure survey. Problems can arise with both components. For example, some countries collect good data in urban centers, but find it much more difficult to collect adequate price data in the countryside, and it is often in the countryside where many poor people live. Urban bias may result in prices that are more relevant for relatively affluent urban

workers than for the poor whose poverty line we are trying to calculate. There are also perennial questions about the representativeness of the markets, and whether the enumerators observe the prices that people actually pay. In India, the villages where prices are collected are updated only very infrequently, if at all. Weights can also be troublesome, especially if not updated sufficiently often. When it was revised in late 1995, the Indian CPIAL (the consumer price index for agricultural laborers) had weights that were 35 years old; it is usually thought that Laspeyres indexes whose weights are held fixed for too long are increasingly likely to overstate the rate of inflation, and thus, in this context, to cause an underestimation of the rate of poverty reduction. (Note, however, that price indexes are *not* at the root of the discrepancy between NAS and survey data; the divergence is in nominal terms.) When updating its poverty lines, the Indian Planning Commission reweights the components of the CPIAL so as to match more closely the purchases of people near the poverty line, using weights from 1973–74, which are still elderly by international standards. (The Bank, in its $1 a day poverty calculations uses a slightly modified version of the CPIAL.) Deaton and Tarozzi (2000) have made independent estimates of price indexes in India which suggest that even the reweighted CPIAL has been rising too quickly. Their alternative estimates therefore show some decline in poverty in the 1990s, though much less than would be the case if the NSS data were scaled up to agree with the NAS data on consumption. I suspect that India is much better provided with good price indexes than are other countries. Indeed, much of the recent debate on the effects of the East Asian crisis on poverty in Indonesia has foundered on the unavailability of good price indexes for rural areas, see Frankenberg, Thomas and Beegle (1999).

Local price indexes are used, not only for updating over time, but also for adjusting poverty

20

lines for urban to rural price differences, as well as for regional differences in prices. (In the World Bank calculations of the number of poor in the world, separate urban and rural indexes are used only for India and for China. In other countries a single index does service for everyone, an expedient which must overstate rural relative to urban poverty.) Again, India is an instructive example. For many years, India had urban and rural poverty lines of Rs 57 and Rs 49 respectively at 1973–74 prices, a differential of 16 percent. Both lines were updated by the implicit price deflator of consumers' expenditure in the national accounts. Since 1993, the Planning Commission has adopted different urban and rural poverty lines for each state, lines that are designed to reflect different prices in different parts of the country, and each of which is updated by a poverty reweighted urban or rural price index specific to that state. For a number of reasons, the implicit urban to rural price difference in the current lines, a difference that was never explicitly measured, has risen to around 40 percent on average, with larger or smaller differentials in individual states. This large urban to rural price difference, for which Deaton and Tarozzi (2000) find no evidence in prices actually paid, yields urban poverty rates higher than rural poverty rates in many states, a finding that seems implausible to most observers and that sits ill with the evidence on other poverty-related indicators. Deaton and Tarozzi recalculate the rates based on their own price indexes, with differentials quite close to the once standard 16 percent, and find that urban poverty is reduced from 32.4 to 18.1 percent in 1993–94, which reduces the total number of Indians in poverty by 23 million people. Price differences within a country can also have major effects on the poverty counts both nationally, as well as in their contribution to the world totals.

2.3 Problems with household surveys

The Bank is to be commended for building so much primary survey data into its poverty counts. This is a great advance over the situation twenty years ago, and the Bank has been the prime mover in bringing these data to bear on the measurement of poverty. But the surveys have a number of limitations for poverty work that need to be kept in mind, some of which might be ameliorated by more work or different approaches. I discuss three of the most important here, (a) coverage, (b) reference periods, and (c) income versus consumption.

Household income and expenditure surveys are good at collecting data on cash that passes through the household, incomes and outlays. They are somewhat less good (but probably still good enough) at collecting data on own production and consumption, However they do not attempt to include in the consumption or income estimates the benefits that people receive from publicly provided goods of one kind or another. Although it would be difficult to collect such information, some relevant data are collected in the surveys (for example, visits to clinics, school attendance, etc.) and perhaps more might be. But if this information is ignored in the poverty counts, we run into problems. In particular, consider the debates about structural adjustment, and the (widely credited) arguments by NGOs and others that the Bank has impoverished the poor by forcing cuts in public expenditure. It is not very useful, as the Bank sometimes does, to argue that there was no impoverishment, because the poverty counts (which take no account of public good provision) are going down. Any program that eliminated government services, closing clinics and schools, and shared the money among the population, would reduce the poverty count as the Bank currently makes it. The Bank has a good deal of experience in benefit incidence analysis, the use of household surveys to allocate the benefits of public expenditures. It would be worth

giving serious consideration to adding some of these estimated benefits to the consumption totals before making the poverty counts.

Different household surveys use different recall periods, even for the same types of goods. Even within the LSMS surveys, there is some variation. Given a true flow of expenditures, the rate of expenditure that is reported is sensitive to the reporting period; longer reporting periods lead to more forgetting, at least for commonplace purchases like food. Again there is a current good example from India. Based on experiments done long ago by Mahalanobis, whose results were quite contrary to other later findings in the literature, the Indian NSS has always used a 30 day recall period for food. Seven or fourteen days is a more common standard in modern surveys, and Scott and Amenuvegbe (1990) have argued that estimates are biased downward after even one day! In the last several years, the NSS, worried by the possible consequences, has done an experiment, with half the households given a seven day reference period for food, and half the (old) thirty day period; there were also some less important changes for goods such as durables and clothing, where a 365 reference period was used, again compared with a 30-day period in the standard survey.

Visaria (2000) has shown that the shorter recall period causes more (food) expenditures to be reported, and that using the official poverty lines and the new data, measured rural poverty in the first six months of 1998 falls from 42.6 percent to 23.6 percent, and urban poverty from 32.9 percent to 20 percent, removing about 175 million people from poverty. (The effect would be somewhat larger for the $1-a-day international line, which is higher than the Indian official poverty line.) Of course, Visaria's calculations offer no apparent help on the issue of differences in trend between the NAS and the NSS versions of consumption. (While it is possible that the

23

relatively under reported foods are those that have been growing more rapidly, preliminary analysis of the data do not support the conjecture.) Somewhat surprisingly, the estimates based on the experimental questionnaire (7-day, 30-day, and 365-day recall, as opposed to 30 days for all goods), show *less* variance across households than do the estimates from the traditional questionnaire, Government of India (2000), Deaton (2000). In consequence, there is no evidence of people being incorrectly characterized as poor simply because they happened to make food purchases outside the shorter recall period.

One question about Visaria's calculations is whether it is legitimate to use the old poverty line with the new reporting periods, and the greater food consumption that they reveal. The answer depends on where one thinks the poverty line comes from and the purpose it serves. Since the Indian line was originally set by the nutritional needs method, perhaps the obvious thing to do it to try to recalculate what the poverty line would have been had the original method been followed using data with the shorted reference period. I shall follow this route below, but it is not the only possible treatment. Instead, one might argue that, although the line was originally set with reference to nutritional norms, the legitimacy and longevity of the line owes more to the general perception that Rs 57 per head in 1973–74 prices is a reasonable poverty line in reference to people's own experiences of earning and spending and their notions of what constitutes poverty. If so, there is no reason to revise it in the light of the new survey data, and Visaria's calculations are correct as they stand. Alternatively, one might argue that the line is acceptable only because people accept the headcount ratio that it generates, in which case the survey data are essentially irrelevant.

If we follow the calorie approach, it turns out that the poverty lines should be revised, not

24

upwards as intuition might suggest, but *downwards*, so that Visaria's calculations if anything

overstate the rate of poverty. The argument is as follows. Pick a level of household per capita

total expenditure, and consider the average of food expenditure over all households at that level,

under the new and old survey methodology. Because the new methodology recovers food

expenditures that were previously forgotten, food expenditures rise on average. But total

expenditure is both defined and measured as the sum of food and non-food expenditures, so that

this too will rise, by the same amount as does food expenditure. (For the purpose of the

argument, ignore the change in survey procedures for non-food.) Figure 1 illustrates. The old (30

day reporting period) Engel curve is the lower curve. For an average household (conditional on

expenditure) at point A, food and total expenditure (or per capita) increase by equal amounts

along the 45-degree line to B, say. Because the slope of the curve is less than 45-degrees—the

marginal propensity to spend on food is less than one—the switch to the 7 day period moves the

Engel curve outwards to the higher curve. The constant level of nutritional needs, represented in

the graph by point N, thus requires a shift downward in the poverty line from the old P to the

new, lower P*. This makes sense. With the new, better reporting period, more food is reported,

so that if we put ourselves back in the shoes of the Planning Commission when the poverty line

was first calculated, their original poverty line would have generated more calories than their

norm, and would therefore have been revised down.

In Deaton (2000), I present evidence on the relevance of Figure 1 using data from the 52[nd]

Round of the NSS gathered during 1995–96. The Engel curves relating per capita expenditure on

food to total household expenditure per capita are in the order predicted in Figure 1, with the 7-

day Engel curve lying on or above the 30-day Engel curve. However, the two curves essentially

coincide below and up to the poverty line so that, unless the ratio of calories to food expenditure is different in the two questionnaires, there is no reason to revise the original poverty line to meet the new questionnaire design.

In summary, Visaria's reduced poverty counts have as much (or more, judged by international survey practice) legitimacy as the official counts, or indeed the $1-a-day calculations. The change in the survey reporting period reduces Indian poverty by as much as the total of number of poor in China! Clearly, the 1.2 billion has a very large margin of error. In the long run, household surveys should be brought into reasonable conformity in their reporting periods. In the short run, the counts should not simply add together counts from surveys with different reporting periods without some sort of correction. This will involve sometimes admittedly arbitrary choices, and perhaps the removal altogether of some surveys. But that is better than adding together numbers that are wildly incommensurate.

The surveys used for the world poverty count sometimes collect data on income, and sometimes on consumption. The latter is the preferred measure (rightly in my view) so that the surveys that collect income (38 out of the 91 in CR, including all the Latin American and Carribean surveys) have to be "converted" to a consumption basis. This is done by scaling down the income data in the survey by the ratio of national consumption to national income (presumably personal disposable income) in the national accounts. This has a certain plausibility to it, but in my judgement, it is unlikely to lead to good results. Here are the arguments:

(a) There is some inconsistency in using the national accounts data in this situation, but not at all when there is a full consumption survey, even when the NAS and survey data are mutually inconsistent. The argument is presumably that there is little choice. Even so, it would seem more

logical to use some average of NAS and survey data to correct the expenditure survey when available, or to make some adjustment to the NAS estimate of the consumption ratio in the case when it is not.

(b) In countries where household surveys collect data on both income and consumption, there is usually little or no evidence of positive saving among poorer households; indeed, the reverse it true. (China seems to be an exception to this generalization.) Even in industrialized countries, survey saving rates are frequently negative until well up the income distribution. One frequently heard argument is that survey-respondents understate income relative to consumption. If so, multiplying incomes by a number that is less than one is a "correction" in the wrong direction! More generally, whatever the accuracy of the consumption and income reports themselves, it remains true that, in most countries, multiplying survey income by the national consumption rate would do a very poor job of replicating survey consumption, certainly at the bottom of the income distribution. In consequence, a proportional correction for saving is not likely to convert income data from a survey into something that looks like the consumption data from a survey. In the usual situation, when the poverty line is below the mode of a unimodal distribution, the resulting underestimation of consumption by the poor (relative to a consumption survey) will tend to overstate poverty rates. From this source alone, my guess is that countries with income surveys tend to be assigned higher poverty rates relative to those with consumption surveys. This contention could be tested for a few countries where surveys collect both kinds of data, or by running a regression across the countries in the data base, and seeing whether those with income surveys have higher poverty, controlling for GDP and other variables. Unfortunately, the degree of heterogeneity across countries is unlikely to permit any simple correction based on the results

27

of such a regression.

(c) Measured consumption is famously less variable than measured income. This is true in logarithms as well as in levels (the variance of the logarithm of consumption is less than the variance of the logarithm of income) so that even if the proportional saving assumption were right on average, the corrected income data would still be too variable relative to measured consumption. Once again, and provided the poverty line is below the mode, consumption poverty will be overestimated by income poverty.

My guess is that it would not be too hard to work out some pattern of corrections that, while hard to defend absolutely, would likely still be better than the current procedures. In the long term, the battle for expenditure surveys will be won, and even the holdouts in Latin America are moving in that direction!

3. Summary, recommendations, and directions for research

The World Bank's worldwide count of the poor starts from a common international poverty line and counts the number of people in each country whose consumption lies below it. The international poverty line, at $1-or $2-a-day is converted into domestic currencies using purchasing power parity exchange rates. I argue that, although the $1-a-day common line has much to recommend it, its dependence on purchasing power parity exchange rates has a number of unfortunate consequences. An arguably better procedure for the future would be to hold fixed (in real terms) the current domestic poverty lines, and not to revise them along with changes in PPP exchange rates induced by updating of base years. Such a program would work best if the current rates were subjected to detailed, local scrutiny, and corrected in a way that would

simultaneously give them local credence without major deviation from the $1-a-day standard. Such corrections will likely be difficult, and must not be interpreted as a license for individual countries to define their own poverty lines with a view to their international political effects. Corrections will be easier if, as I suspect, most anomalies are the result of inappropriate PPP exchange rates, perhaps because interpolation from other countries gives a poor answer, or because the bundle of goods being priced is wildly inappropriate for the poor.

Other problems relate less to international comparability than to the calculation of poverty within each country. The proliferation of household surveys, for which the Bank should take much credit, has greatly enriched the basis for good poverty calculations around the world. But the very availability of expenditure surveys has highlighted an old problem that seems to have recently become much worse, which is a divergence between estimates of average consumption based on surveys and on national income accounts. The Bank is probably the only organization that is capable of leading an international research initiative to address these discrepancies. Until it does so, policymakers and policy commentators will be invited to choose their evidence according to their political predilections. Furthermore, neither the Bank nor anyone else will be able to make any well-supported statement about what is surely a central issue in economic development, whether growth has reduced poverty in the past, or is likely to do so in the future.

There is also much scope for the improvement of survey practice and the convergence of household expenditure surveys towards international best practice standards. The experimentation with recall periods by the Indian National Sample Survey, although it has sown some confusion in the short run, is an excellent example of the sort of methodological work that is needed, and which statistical offices are usually reluctant to perform. If nothing else, these

experiments have highlighted the extreme sensitivity (175 million people less in poverty, the headcount ratio cut by a half) of poverty measurement to what many would previously have regarded as a technical, if not downright esoteric issue of survey design. Here too is an important research program; we need to understand much better than we do now how the measurement of consumption is affected by the design of consumer expenditure surveys—not only reference periods, but such issues as the detail of the questionnaire, the need for bounding interviews and repeat visits, multiple versus single respondents, whether ownership information on durables is worth collecting, and so forth. Although the Bank's Living Standards Measurement Study has done much to normalize the use of enriched expenditure surveys in many countries, it has done little experimentation and has not greatly contributed to the settling of design questions, see Deaton and Grosh (2000).

Another important research issue is finding better ways to set the poverty line. Within the standard realm of income (consumption) poverty, there is much to be said for giving greater weights to people's own reports of consumption adequacy. Ravallion has done sterling work in this direction, Ravallion (1998) and Ravallion and Pradhan (2000), and there is widespread interest around the world in these methodologies. In the U.S. for example, Gallup poll questions people how much would it take for a family like theirs to 'just get by' in your community yield sensible poverty lines, at least once outliers have been removed. People can also be asked to self-assess their poverty, as in Ravallion's work, and if such schemes could be put on a sound footing, they raise the tantalizing prospect of reducing our dependence on extremely expensive, time-consuming, and most likely inaccurate consumption surveys. One or two questions are a good deal cheaper to ask than two or three hundred! Recent work in South Africa shows strong

correlations between measures of financial, physical, and mental health, Case and Wilson (2000).

But there are real concerns about "adaptation," that people do not perceive themselves to be better off even in the face of large increases in real income, essentially because they adapt their expectations to their circumstances, as in Easterlin's (1995) famous finding that subjective levels of happiness did not increase among Japanese consumers between 1958 and 1987 in spite of a fivefold increase in real per capita income. The Bank might find itself successfully promoting growth with a concomitant reduction in income poverty, but then being negatively assessed because people do not report themselves to be better-off. The Washington consensus derailed by the hedonic treadmill!

All of us who have been exposed to the field experience have been impressed by the prominence of health concerns in what people tell you about their poverty. Income, housing, and jobs tend to predominate when health is normal, but if someone get sick, is hit by a car, or has a friend or relative who is raped or murdered, income poverty recedes into the background in people's perceptions. Many millions of people around the world are going to die from AIDS, with untold misery and deprivation. It would be a terrible thing if the Bank dealt with this only in terms of its effects on income poverty, which it already shows some signs of doing. (For example, the July 2000 PovertyNet Newsletter writes that "A decade ago, HIV/AIDS was regarded primarily as a serious health crisis. Today, it is clear that AIDS is a development crisis." If per capita income were to rise as people died, would that make the impact of AIDS somehow less severe?) My own view is that the Bank should be backing away from its current too concentrated focus on the income headcount numbers, and that the emphasis should move to a much wider range of other measures, focusing on a wider range of deprivations that are as, or

more, important than deprivation of income.

References:

Bardhan, Pranab K., 1974, "The pattern of income distribution in India: a review," in Srinivasan and Bardhan, 103–38.

Bhalla, Surjit S., 2000a, "Growth and poverty in India: myth and reality," New Delhi, processed. (June).

Bhalla, Surjit S., 2000b, "Trends in world poverty: ideology and research," New Delhi, processed. (June).

Case, Anne and Francis Wilson, 2000, "Health and wellbeing in South Africa: evidence from the Langeberg survey," Princeton, NJ. Research Program in Development Studies, processed. (December)

Chen, Shaohua, and Martin Ravallion, 2000, "How did the world's poorest fare in the 1990s?" Washington, D.C., The World Bank, processed.

Deaton, Angus, 2000, "Preliminary notes on reporting periods in the Indian NSS, 52nd through 54th Rounds," Princeton, N J. Research Program in Development Studies.(October).

Deaton, Angus and Margaret Grosh, 2000, "Consumption," Chapter 5 in Margaret Grosh and Paul Glewwe, eds., *Designing household survey questionnaires for developing countries: lessons from 15 years of the Living Standards Measurement Study*, Volume 1, Washington, DC. The World Bank.

Deaton, Angus, and Alessandro Tarozzi, 2000, "Prices and poverty in India," Princeton, July 2000.

Easterlin, Richard A., 1995, "Will raising the incomes of all increase the happiness of all?" *Journal of Economic Behavior and Organization*, **27**, 35–47.

Frankenberg, Elizabeth, Duncan Thomas, and E. Beegle, 1999, "The real costs of Indonesia's economic crisis," Santa Monica, Ca. RAND Labor and Population Program Working Paper 99–04.

Filmer, Deon and Lant Pritchett, 1998, "Estimating wealth effects without expenditure data-or tears: with an application to educational enrollments in states of India," Washington, D.C., The World Bank, Policy Research Working Paper 1994, processed. (October.)

Government of India, 1993, *Report of the expert group on the estimation of the proportion and number of poor*, Delhi. Planning Commission.

Government of India, 2000, *Choice of reference period for consumption data*, National Sample Survey Organization.

Groves, Robert M., and Mick P. Couper, 1998, *Nonresponse in household surveys*, New York. Wiley.

Mukherjee, M., and G. S. Chaterjee, 1974, "On the validity of NSS estimates of consumption expenditure," in Srinivasan and Bardhan, 139–47.

National Research Council, 1995, *Measuring poverty: a new approach*, Washington, D.C. National Academies Press.

Pradhan, Menno and Martin Ravallion, 2000, "Measuring poverty using qualitative perceptions of welfare," *Review of Economics and Statistics*, August, forthcoming.

Ravallion, Martin, 1998, "Poverty lines in theory and practice," LSMS Working Paper No. 133, Washington, D.C., The World Bank.

Ravallion, Martin, 2000a, "Should poverty measures be anchored to the national accounts?" Washington DC. The World Bank. processed. (July).

Ravallion, Martin, 2000b, "Do national accounts provide unbiased estimates of survey based estimates of living standards?" Washington, DC. The World Bank, processed. (August.)

Ravallion, Martin and Benu Bidani, 1994, "How robust is a poverty profile?" *World Bank Economic Review*, 8(1), 75–102.

Ravallion, Martin and Binayak Sen, 1996, "Monitoring poverty in Bangladesh," *Economic Development and Cultural Change*, 44(4), 761–92.

Scott, Christopher, and Ben Amenuvegbe, 1990, "Effect of recall duration on reporting of household expenditures: an experimental study in Ghana," Washington, D.C. The World Bank, Social Dimensions of Adjustment in Sub-Saharan Africa Working Paper No. 6.

Srinivasan, T. N., 2000, "Growth and poverty alleviation: lessons from development experience," New Haven, CT. Department of Economics, Yale University, processed. (Nov).

Srinivasan, T. N., and P. K. Bardhan, 1974, *Poverty and income distribution in India*, Calcutta, Statistical Publishing Society.

Srinivasan, T. N., P. N. Radhakrishnan, and A. Vaidyanathan, 1974, "Data on distribution of consumption expenditure in India: an evaluation," in Srinivasan and Bardhan, 148–162.

Triplett, Jack E., 1997, "Measuring consumption: the post-1973 slowdown and the research issues," *Federal Researve Bank of St Louis Review*, May/June, pp 9–42.

Visaria, Pravin, 2000, "Poverty in India during 1994–98: alternative estimates," Institute for Economic Growth, New Delhi, processed. (June 9)
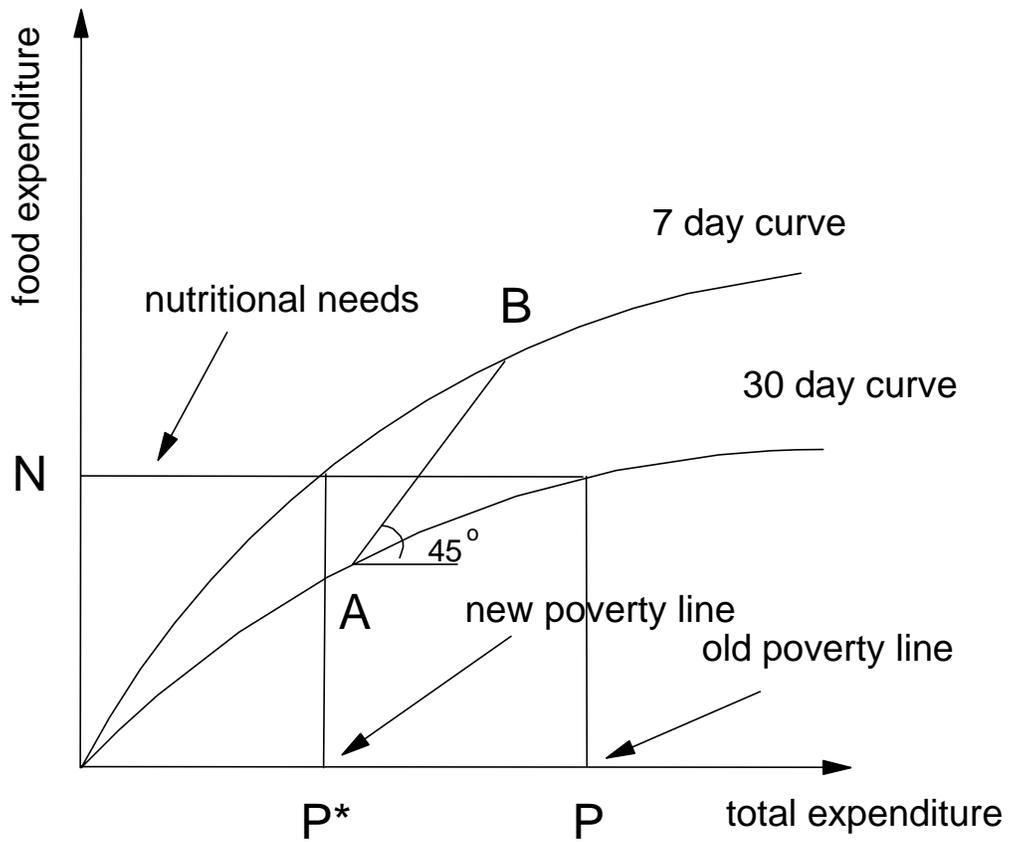
Figure 1: **The effects of a shorter reporting period on a nutritional poverty line.**