Well-being Measurement and Analysis Technical Notes

TN 1 Measuring poverty and analyzing changes in poverty over time

This note provides mathematical expressions for the most commonly used poverty measures and for their decomposition by sector or, more generally, by group. The note focuses on the first three poverty measures of the so-called FGT class (Foster, Greer, and Thorbecke 1984), namely the headcount, the poverty gap, and the squared poverty gap.

Poverty measures

Poverty Headcount: This is the share of the population which is poor, i.e. the proportion of the population for whom consumption or income *y* is less than the poverty line *z*. Suppose we have a population of size *n* in which *q* people are poor. Then the headcount index is defined as:

$$H = \frac{q}{n}$$

Poverty Gap: The poverty gap, which is often considered as representing the depth of poverty, is the mean distance separating the population from the poverty line, with the non-poor being given a distance of zero. The poverty gap is a measure of the poverty deficit of the entire population, where the notion of "poverty deficit" captures the resources that would be needed to lift all the poor out of poverty through perfectly targeted cash transfers. It is defined as follows:

$$PG = \frac{1}{n}\sum_{i=1}^{q}\left[\frac{z - y_i}{z}\right]$$

where y_i is the income of individual *i*, and the sum is taken only on those individuals who are poor (in practice, we often work with household rather than individual income, but individual income can still be defined as being equal, say, to the per capita income of the household). The poverty gap can be written as being equal to the product of the income gap ratio and the headcount index of poverty, where the income gap ratio is itself defined as:

$$PG = I * H, \text{ with}$$

$$I = \frac{z - y_q}{z} \quad \text{where} \quad y_q = \frac{1}{q}\sum_{i=1}^{q} y_i \quad \text{is the average income of the poor.}$$

It must be emphasized that the income gap ratio *I* in itself is not a good measure of poverty. Assume that some households or individuals who are poor but close to the poverty line are improving their standards of living over time, and thereby become non-poor. The Income gap ratio will increase because the mean distance separating the poor from the poverty line will increase (this happens because some of those who were less poor have emerged from poverty – so that those still in poverty are on average further away from the poverty line), suggesting a

deterioration in welfare, while nobody is worst off and some people are actually better off. Although the income gap ratio *I* will increase, the poverty gap itself *PG* will decrease, because the headcount index of poverty will decrease, suggesting an improvement towards poverty reduction. The problem with the income gap ratio is that it is defined only on the population that is poor, while the poverty gap is defined over the population as a whole.

As mentioned above, the poverty gap is a useful statistics to assess how much resources would be needed to eradicate poverty through cash transfers **perfectly targeted** to the poor. Assume for example that the poverty gap is equal to 0.20. This means that the cash transfer needed to lift the poor out of poverty each poor person represents 20 percent of the poverty line. If the mean income in the country is equal to twice the poverty line, the cash transfer would represent 10 percent of the country's mean income. Now, if it is the mean income of the non-poor which is equal to twice the poverty line, and if half the population is poor, it can be shown that the tax rate that would have to be imposed on the non-poor to lift the poor out of poverty with perfectly targeted transfers would be 20 percent again. If the mean income of the non-poor is equal to four times the poverty line, under the same assumption the necessary tax rate would be 10 percent. Such simple simulations can be used to communicate in an intuitive manner the meaning of the poverty gap. In practice however, given that perfectly targeted cash transfers to eradicate poverty are neither feasible nor necessarily a good thing (high tax rates could stifle economic growth and thereby future poverty reduction), one must be careful in their use.

Squared Poverty Gap: This is often described as a measure of the severity of poverty. While the poverty gap takes into account the distance separating the poor from the poverty line, the squared poverty gap takes the square of that distance into account. When using the squared poverty gap, the poverty gap is weighted by itself, so as to give more weight to the very poor. Said differently, the squared poverty gap takes into account the inequality among the poor. It is obtained as follows:

$$P2 = \frac{1}{n}\sum_{i=1}^{q}\left[\frac{z - y_i}{z}\right]^2$$

The headcount, the poverty gap, and the squared poverty gap are the first three measures of the Foster-Greer-Thorbecke class of poverty measures. The general formula for this class of poverty measures depends on a parameter α which takes a value of zero for the headcount, one for the poverty gap, and two for the squared poverty gap in the following expression:

$$P\alpha = \frac{1}{n}\sum_{i=1}^{q}\left[\frac{z - y_i}{z}\right]^{\alpha}$$

It is important to use the poverty gap or the squared poverty gap in addition to the headcount for evaluation purposes, since these measure different aspects of income poverty. Indeed, the basing evaluation on the headcount ratio would consider as more effective policies which lift the richest of the poor (those close to the line) out of poverty. On the basis of the poverty gap *PG* and the squared poverty gap *P2,* on the other hand, puts the emphasis on helping those who are further away from the line, the poorest of the poor.

Decompositions for changes in poverty over time

Two main decompositions have been used in the literature to analyze changes in poverty over time. The first decomposition deals with shifts in poverty between sectors or groups (Ravallion

and Huppi 1991). The second decomposition deals with the contribution of income growth and changes in inequality to changes in poverty (Datt and Ravallion 1992; Kakwani 1997).

Sectoral decomposition

The poverty measures of the FGT class are additive. This means that the poverty measure for the population as a whole is equal to the weighted sum of the poverty measures for the population subgroups, with the weights defined by the population shares of the subgroups. This additive property makes it feasible to analyze the contribution of various population subgroups to changes in overall poverty over time. Assume that households or individuals can be classified according to various sectors in the economy. These may be industrial sectors, geographic sectors (urban versus rural), or any other sectors that the analyst may suggest. The overall change in poverty over time can be decomposed into: 1) changes in poverty within specific sectors, or intra-sectoral changes, 2) changes in poverty due to changes in the population shares of sectors, or inter-sectoral changes, and 3) changes due to the possible correlation between intra-sectoral and inter-sectoral changes, or interaction effect. Denote by P_{it} the poverty measure in sector i at time t; there are m sectors ($i=1,\ldots, m$), with population share n_i in sector i, and two periods (1 and 2). Then, the overall change in poverty is equal to:

$$\Delta P = \underbrace{\sum_{i=1}^{m} n_{i1}(P_{i2} - P_{i1})}_{\text{Intra-sectoral}} + \underbrace{\sum_{i=1}^{m} P_{i1}(n_{i2} - n_{i1})}_{\text{Inter-sectoral}} + \underbrace{\sum_{i=1}^{m} (P_{i2} - P_{i1})(n_{i2} - n_{i1})}_{\text{Interaction effect}}$$

Growth and inequality decomposition

Changes in poverty rates can also be decomposed into changes due to economic growth (or mean income) in the absence of changes in inequality (or income distribution), and changes in inequality in the absence of growth. Denoting by $P(\mu_t, L_t)$ the poverty measure corresponding to a mean income in period t of μ_t and a Lorenz curve L_t, the decomposition is:

$$\Delta P = \underbrace{[P(\mu_2, L_r) - P(\mu_i, L_r)]}_{\text{Growth impact}} + \underbrace{[P(\mu_r, L_2) - P(\mu_r, L_1)]}_{\text{Inequality impact}} + \underbrace{R_r}_{\text{Residual}}$$

The first component is the change in poverty that would have been observed if the Lorenz curve had remained unchanged, while the second component is the change that would have been observed if mean income had not changed. The last component is a residual.

TN 2 Estimating poverty lines: The example of Bangladesh

As noted in the main text, to measure poverty, one needs: a) an indicator of well-being or welfare such as per capita caloric intake or per capita expenditure; b) a threshold (the poverty line) to which each individual or household's welfare can be compared; and c) a poverty measure. Differences in poverty estimates can result from differences in the choice of the indicator, the threshold, or the poverty measure. Using an example from Bangladesh, this note focuses note on differences in the choice of the indicator and the threshold. Specifically, the focus is on: a) the degree of representativeness of the indicator, by which it is meant the extent to which the chosen indicator of welfare is able to capture the well-being of the households at a

sufficiently broad level; and b) the degree of consistency of the threshold (poverty line), by which it is meant the extent to which the chosen threshold represents similar levels of well-being over time and across groups, ensuring that the poverty estimates can be used for valid comparisons of poverty over time and across groups.

Three main methods have been used for estimating poverty in Bangladesh. The methods differ in terms of their indicator of welfare and their approach to the threshold or poverty line: direct caloric intake, food energy intake, and cost of basic needs, as summarized in Table 26. The drawback of the direct caloric intake method is that its indicator is not representative, while the drawback of the food energy intake method is the lack of consistency of its threshold or poverty line. The cost of basic needs method can be deemed to be both representative and consistent for comparisons over time and across groups, at least more so than the alternative methods.

Table 26. Strengths and Weaknesses of Alternative Methods for Poverty Measurement

	Direct Caloric Intake	**Food Energy Intake**	**Cost of Basic Needs**
Indicator	Caloric intake	Expenditure (or income)	Expenditure (or income)
Threshold	2,122 kcal/person-day	Expenditure level at which household members are expected to reach caloric intake threshold	Expenditure level at which household members are expected to meet basic needs (food and non food)
Measure	Headcount or other	Headcount or other	Headcount or other
Strengths and Weaknesses	Indicator not representative Threshold consistent (for monitoring caloric intake)	Indicator representative Threshold not consistent (for real expenditures)	Indicator representative Threshold consistent (for real expenditures)

The Bangladesh Bureau of Statistics (BBS) has first relied on the direct caloric intake method to measure poverty, which considers as poor any household not meeting the nutritional requirement of 2,122 kcal per day and per person. That is, on the basis of the quantity of food consumed by each household, the BBS computed per capita caloric intakes and considered as poor any household whose intake was below 2,122 kcal per person on a daily basis. The difficulty with this method is that it equates poverty with malnutrition or a proxy thereof. If we consider poverty as a lack of command on basic goods and services, measuring poverty by caloric intake only is unlikely to represent adequately the state of deprivation of the poor.

The BBS has also relied on the food energy intake method to compute poverty lines. The idea with this method is to find the value of per capita consumption at which a household can be expected to fulfill its caloric requirement. That is, the poverty line is defined by the level of per capita consumption at which people can be expected to meet this requirement. This represents a methodological improvement in terms of representativeness because the food energy intake method provides a monetary rather than purely nutritional concept of poverty. Yet, the method suffers from major deficiencies in terms of consistency in that the poverty lines generated might not represent an identical level of welfare (specifically, an identical purchasing power in real terms) over time or across groups, and hence poverty comparisons may not be valid.

To illustrate the weaknesses of the food energy intake method, imagine that following a decrease in their income, households change their consumption patterns to consume less expensive food. Although their level of welfare has decreased (and hence poverty must have increased), the researcher using the food energy intake method may well find that poverty has decreased. This is because in buying lower quality food, households meet their caloric

requirement at a lower expected level of per capita consumption expenditure, and hence the estimated poverty line will be lower. This scenario is not hypothetical: in fact, the poverty line in urban areas computed by the BBS with the food energy intake method in 1988-89 was lower than that computed in 1985-86, although prices of most consumption items increased between the two years. Surely, the lower BBS poverty lines of 1988-89 did not represent the same standard of living as the poverty line of 1985-86.

Recently, the BBS has adopted the cost of basic needs method for measuring poverty. With this method, an absolute poverty line is defined as the value of consumption needed to satisfy minimum subsistence needs. Difficulties arise in specifying these needs as well as the most appropriate way of attaining them. For food consumption, nutritional requirements can be used as a guide, as is the case with the other methods. In practice, this is often restricted to calorie (and possibly protein) requirements, but even then there is the question of which food basket to choose in order to meet the requirements. Specifying minimum requirements for nonfood consumption is more difficult, and various methods have been proposed for dealing with non-food basic needs. Another issue relates to the adjustments that must be made for differences in the cost of food and non-food items between regions, and possibly over time, either when the survey has been carried over a relatively long period of time such as one year, or when one is using several surveys for poverty monitoring.

The first step consists in the definition of a bundle of food items meeting a given nutritional requirement (2122 kcal per person and per day). Many food bundles can provide this requirement. The bundle used includes rice, what, pulses, milk, mustard oil, beef, fresh water fish, potatoes, other vegetables, sugar, and bananas. It could be argued that using a common food bundle for the whole country is inadequate because consumption patterns of households may vary across areas (e.g. households in coastal areas may eat more fish). Households in different regions might substitute some goods for others if prices vary by area. Other difficulties such as seasonality in food prices, potential omitted variables or selectivity bias in the choice of food items consumed, or errors of measurement in the database for the imputation of food produced and consumed at home, may lead to bias in the estimates of food prices. In the case of Bangladesh, these considerations are considered not to be too problematic, but in other countries, adjustments may have to be made to the food bundle in various areas.

The second step consists in estimating the cost of the food bundle. To do so prices by geographical area are computed for each component of the food bundle. There are different methods to compute these prices. A first method consists in taking the average prices paid by households in each region (where, for a given household, the prices of the various food items in the food bundle are obtained by dividing the reported food expenditure by the quantity consumed). Because the poor tend to buy goods of lower quality, they usually face on average lower prices than the non-poor. Taking the regional average over all households therefore tends to overestimate prices faced, and to result in a higher poverty line and higher poverty measures. In addition, imagine a situation in which the welfare of the non-poor increase over time while that of the non-poor does not change. Because the non-poor will tend to buy goods of higher quality at higher prices, the poverty line will be higher, suggesting an increase in poverty. A second method consists in taking the average prices over the poor population only. A third method, used by the BBS, consists in using regressions to estimate the differences between regions in the prices paid by households for their food. Having estimated the cost of each food item j in each region k, and denoting these prices by P_{jk}, food poverty lines can be computed for each region k as $Z_{kf} = \Sigma_j P_{jk}F_j$, where F_j is the per capita quantity of food item j in the basic food bundle.

Once the food component of the poverty lines has been estimated, the third step consists in the estimation of a reasonable allowance for non food consumption. Various methods can be used (Ravallion 1994). The BBS considered two methods. The first method is based on the amount of non-food expenditures for the households (in geographical area k) whose *total* consumption is equal to the regional food poverty line Z_{kf} is computed. The non-food expenditures of the households in this first case must be necessities since the households are giving up food expenditures which are considered as necessary in order to buy non-food items. In the second method, the share of the non-food expenditures for the households whose *food* expenditure is equal to the food poverty line was computed. On average, under the second case, the households "near the poverty line" are able to meet their nutritional requirements, and it the non-food allowance computed this way will be more generous than that in the first case. Various techniques (both parametric and non-parametric) can be used to estimate the non-food components of the regional poverty lines, but the details go beyond the scope of this note. The two estimates, specific to each region, can be denoted by Z^L_{kn} and Z^U_{kn} (k is the region, n stands for non-food, and the subscript L and U refer to the "Lower" and "Upper" allowances for non-food consumption obtained respectively from the first and second cases above). Then, the two overall poverty lines which include provisions for both food and non-food basic needs are defined as $Z^L_k = Z_{kf} + Z^L_{kn}$ and $Z^U_k = Z_{kf} + Z^U_{kn}$, with Z^L_k being smaller than Z^U_k.

As mentioned in the main text, once computed, the regional and time adjustments can be used in two ways. Either one adjusts the indicator of consumption or income with a price deflator for each region and period and compare that indicator to a unique reference poverty line, or one keeps a different poverty line for each region or period to which the indicator of income or consumption is compared. The final results in terms of poverty comparisons will be the same since these are simply two alternative ways to use the same adjustments in the poverty comparisons between regions and household groups, or over time.

Figure 7 provides the results of the estimation of the poverty lines for the various regions for one of the survey years. The vertical axis provides the level of the poverty line in Taka per person and per month. The horizontal axis represents 14 different regions in the country. The food poverty line is the lowest one. The middle line represents the total poverty line with the "lower" non-food allowance. The top line represents the total poverty line with the "upper" allowance for non-food consumption. Clearly, there are large differences in the cost of living between areas. Not surprisingly, the area #1, which represent the capital city of Dhaka, has the highest cost of living. Overall, there are large variations by city and/or area within urban and rural areas.

Figure 7. Food, lower and upper poverty lines by area

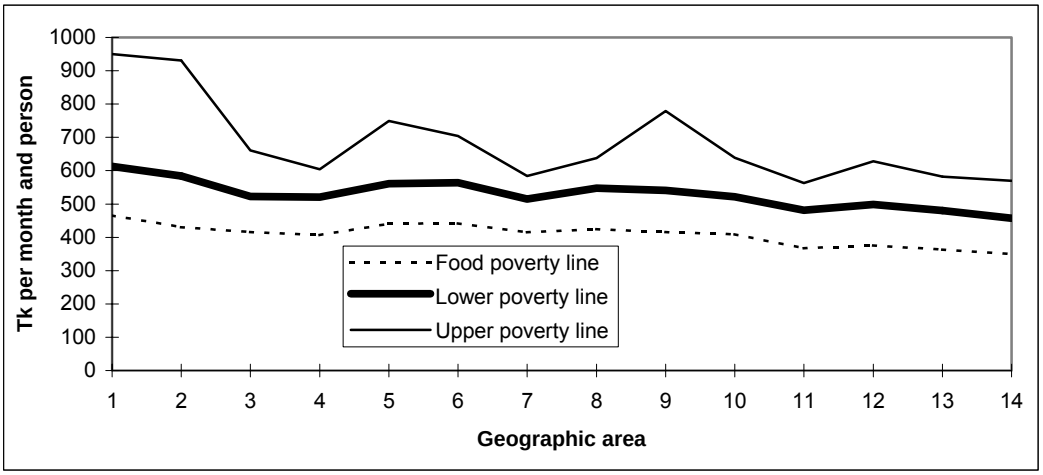

Source: Wodon 1997a

TN 3 Estimating the indicator of well-being: The example of consumption in Uganda

This note presents an example of the types of adjustments that one may have to make for the indicator of well-being (income or consumption) computed at the household level in order to better reflects the underlying level of welfare that one is trying to capture, and to make appropriate comparisons over time and between regions. The illustration is based on the Ugandan experience. Table 27 reports the estimates of consumption per capita as calculated in the official survey reports (i.e., before adjustments) as well as after adjustments. The adjustments made by the authors fall into three categories: adjustments for sampling design, for questionnaire design and for prices.

Table 27: Mean consumption per capita before and after adjustments (Ush. Per month)

	IHS 92/93	MS-1 93/94	MS-2 94/95	MS-3 95/96	MS-4 97/98
Before adjustments (official reports)	11574	13195	15221	17499	20540
1- Adjustment for geographic coverage	11786	13501	15388	17721	20747
2- Adjustment for public transport fares	11981	-	-	-	-
3- Adjustment for food consumed at home	12769	14748	16643	18568	21976
4- Adjustment for differences in regional prices	13187	15267	17064	18973	22139
5- Adjustment for inflation (1989 prices)	5452	5825	6058	6187	6353
6- Adjustment for re-weighting MS-1	5452	5718	6058	6187	6353

Source: Appleton et al. (1999)

1. **Sampling:** The authors used two surveys – the Integrated Household Survey (IHS) and the Monitoring Surveys (MS), both of which have a sampling frame based on the 1991 census and large samples (10000 households in the IHS, 5000 in each MS).

- Changes in geographic coverage.. Security problems led to the exclusion of a few districts (Kitgum, Gulu, Kasese and Bundibugyo) from MS-4. To ensure comparability over time, the authors excluded the four districts from all calculations. These districts are relatively poor so that their omission raises mean consumption per capita by 1.8 per cent in the IHS and 2.3 per cent in the MS-1. (Adjustment number 1).

- Seasonality: The IHS was conducted throughout an entire year but MS-1 and MS-2 were conducted only during certain months of the year. This is problematic since food consumption is reported for a short recall period only and subject to seasonal variations. The authors chose not to adjust for seasonality because of data limitations, and because of an assessment that an adjustment would not affect the main conclusions of the analysis.
- Panel data: The MS surveys have a panel element, with half the enumeration areas being revisited and within those areas, half of the households being revisited. This led to a very high rate of attrition among households and the panel structure was abandoned in MS-4. The authors compared the mean levels of per capita consumption in MS-2 and MS-3 amongst panel and non-panel households, and concluded that no adjustment was needed for attrition. However, after the Ugandan Statistics Department changed its system for weighting households to take into account the panel structure of the data, because this had not been done for the MS-1 survey, the authors amended the weights for this survey, which did alter consumption estimates, raising them in rural areas and lowering them in urban. (Adjustment number 6)

2. **Questionnaire design:** The MSs have similar questionnaires on consumption, while the IHS has more items. The way in which the data are recorded differ in both surveys, both in terms of what the interviewer must do and in terms of the recall periods. In addition, the HIS has information on health and education expenditures at the individual level while the MSs provide this information at the household level. To gauge the effects of these differences, the actual composition of expenditures was compared across surveys. The expenditures were similar across the IHS and MSs, but there was a discrepancy in the share of expenditures on transport and communications reflecting a printing error. To adjust for this, the authors imputed expenditures for transportation using regional shares in the MS-1. Omission from the IHS of health expenditures for one district was dealt with in a similar manner. These adjustments raised the mean consumption figure for the IHS by 1.7%. (Adjustment number 2)

3. **Prices:** Three adjustments were made to obtain consumption estimates in constant prices.

- Valuation of home food consumption: Home food consumption expenditures were revalued to be at market prices. This was done using median unit values from the surveys separately for urban and rural areas in each of four regions (eight sets of prices were computed). The adjustment increased the value of home food consumption by around 30%. (Adjustment 3).
- Regional variation in food prices: Food prices are markedly higher in some areas, particularly urban areas, than others. Median unit values for purchases of major food items were used to construct regional food price indices for each survey. Non-food prices, however, were assumed to be constant across the country. This may be problematic since, as suggested in the previous technical note (Figure 7), variations in non-food prices between regions (and over time) are typically even larger than variations in food prices. (Adjustment 4).
- Inflation over time and within surveys: The composite national CPI was used as price deflator for expenditures recorded in the MS surveys, converting the data into 1989 prices. During the implementation of the IHS, prices increased substantially, so that at the end of the sampling period they were 30% higher than at the beginning. To take this into account, the IHS expenditures we deflated using monthly rather than yearly CPI data. (Adjustment 5)

TN 4 Poverty Maps and their use for targeting

A poverty map is a geographical profile of poverty, indicating in which parts of a country poverty is concentrated. Such maps can play an important role in guiding the allocation of public spending to reduce poverty. A poverty map is most useful if it can be constructed at a fine level of geographical disaggregation. Unfortunately, accurate geographical disaggregration requires work with large data sets such as population censuses which typically do not contain detailed income or spending information because collecting such data for the whole country is very expensive.

There are two ways of dealing with this problem. First, one can develop poverty maps on the basis of indices of welfare constructed by combining information on variables such as access to water, electricity, sanitation, or the education levels of the household head.

A second method consists in combining survey and census data. The idea is to develop simple models in which consumption or income are a function of such factors as housing, employment, household characteristics (such as size and composition), and education variables (see **Technical Note 8**). The estimated parameters are then used in the census to predict household consumption. While consumption is not in the original census data set, it is artificially added by using structural relationships derived from the survey. This approach can be pursued only if three data requirements are met. First, a household survey with consumption and other household characteristics must be available, and should correspond roughly to the same period covered by the census. Second, unit record level census data must be available for analysis. Third, a sufficient number of variables that are used to predict consumption must be available in the survey **and** in the census—otherwise it would not be possible to use the models based on the surveys to predict census consumption data.

The first step is to estimate a model of consumption or income using household survey data. As mentioned, the only variables that can be used to predict consumption or income are variables that are also available in the census. The second step is to apply the parameter estimates from the regressions to the census data. It is then possible to derive consumption or income estimates based on individual household characteristics in the census. These consumption or income data then allow the estimation of the probability of each household in the census data set of being poor. Since the estimates are derived from an imperfect model, they will have a margin of error attached to them, and this has to be taken into account (see Section 2.2.4 in the main text). While the accuracy of each household-level estimate may be low, at a more aggregate level, these errors will tend to compensate each other and regional or district estimates will be relatively accurate.

Ultimately, the optimal degree of disaggregation will depend on many factors. First, it will depend on the purpose for which the poverty map is built. Is it, for example, intended to identify government administrative areas so that the desired level of disaggregation is the same as the level of local government? Or is it intended to identify poor villages or neighborhoods within an administrative area so that community-level project interventions can be better targeted? Second, the choice of disaggregation level will depend on the extent to which the parameter estimates from a regression estimated at the regional level can be assumed to apply to sub-regional breakdowns. Indeed, the parameters have to be estimated at the level for which the household survey is representative (usually broad regions). Applying the parameters to all households in the census who belong to that region relied on the assumption that that, within a region, the model of consumption or income is the same for all households regardless of which community they live in. Third, the desired degree of disaggregation will also depend on the availability of other sources of information on the poverty of individuals that might be available locally.

Probably the most useful practical application to which this methodology can be devoted lies in combining consumption-based poverty maps with other indicators of well-being, opportunity, and access for sectoral investment planning. For example, a map documenting, say, regional patterns of access to primary health care centers can be overlaid against a consumption- or income-based poverty map. The map can use other poverty statistics, such as the poverty gap. It might help policymakers decide where to prioritize efforts to expand access to primary health centers; health investment planning can prioritize the poorest areas with the lowest health coverage. Furthermore, a close correlation between, say, regional patterns of rural poverty and road access might also offer clues on possible causes of poverty. This type of exercise could be undertaken for a wide range of indicators: levels of health and education; ethnicity; access to infrastructure and other public services; or land quality and ecology

See http://www.worldbank.org/poverty/inequal/povmap/index.htm for references on poverty maps and Hentschel et. al. (2000) for an application.

TN 5 Stochastic Dominance Tests

When comparing poverty measures over time or between groups, it is important to test the robustness of the observed changes in poverty indices. Indeed, the observed changes might depend on the selected poverty line and at the extreme, using two different poverty lines can suggest changes in opposite directions. Comparing poverty measures using stochastic dominance techniques can help in establishing the robustness of ordinal poverty rankings.

First-order statistical dominance involves comparing the cumulative distribution functions for the indicator of well-being (income or consumption) for each of the survey years, or for the various groups of households for which poverty comparisons are made. One distribution ‘dominates’ another if the income distribution function for that year or that household group lies above that of the year or other group at all levels of income or consumption. If one finds that first-order dominance holds between two different years, or between two different groups, this implies that all FGT poverty measures – including the headcount, the poverty gap, and the squared poverty gap — in the first year or group are higher than in the other year or group for all poverty lines.

Second-order dominance tests involve analyzing 'deficit' curves, or integrals of the cumulative income distribution functions, and similarly determine whether poverty has improved or worsened over time for all poverty measures of the order of the poverty gap or higher, such as the squared poverty gap. Still higher levels of dominance can be established, and multivariate stochastic dominance can be used in the context of multidimensional distributions of poverty.

Figure 8 shows an example from Ghana, comparing the cumulative distribution functions (for first order stochastic dominance) of 1991/92 and 1998. The cumulative distribution for 1998 always lies below that for 1991/92, which implies that poverty has unambiguously fallen over this time period. If the two distributions had crossed, then poverty measures would have shown an increase for all the poverty lines which are in the range for which the second distribution is below the first, and a decrease for the others. See also Atkinson and Bourguignon (1982).

Figure 8: Cumulative Distribution Functions to Establish First-Order Dominance

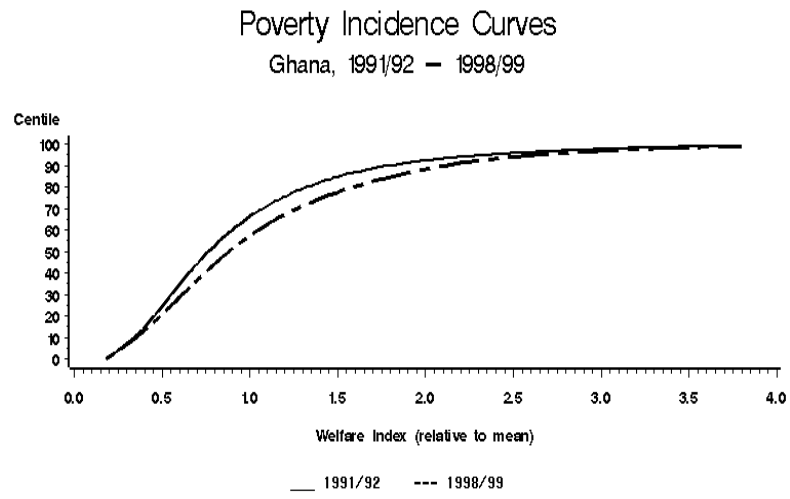

So-called sequential stochastic dominance tests have also been developed to check for the robustness of poverty comparisons to assumptions regarding differences in household needs according to household size apart from assumptions regarding the poverty lines (for recent applications, see for example Duclos and Makdissi, 2000; Makdissi and Wodon, 2001).

TN 6 Applying poverty measurement tools to non-monetary indicators

Although poverty has been traditionally measured in monetary terms, it has many other dimensions, such as lack of access to public services related to education, health, and infrastructure, to deficient social relations, to insecurity and vulnerability, and to low self-confidence and powerlessness. In some cases, it is feasible to apply the tools which have been developed for poverty measurement to non-monetary indicators of well-being. The requirement for being able to apply the tools of poverty measurement to non-monetary indicators is that it must be feasible to compare the value of the non-monetary indicator for a given individual or household to a threshold or "poverty line" under which it can be said that the individual or household is not able to meet its basic needs. The idea can be illustrated with three examples:

- Health and nutrition poverty: Morris, Flores, and Zúniga (2000) analyze the potential for targeting nutrition intervention programs in Honduras. The "nutrition poor" are defined as stunted children, i.e. children who have a measure of height for age at least two standard deviations below international standards. The authors simulate a nutrition intervention reaching 20 percent of the children with a gain of half a standard deviation for beneficiaries, with different approaches for targeting the beneficiary children. The impact of the intervention is obtained by computing the incidence of stunting (the headcount for the nutrition poor) as well as the malnutrition gap and the quadratic malnutrition gap before and after the intervention. They find that the simulated program has the potential to substantially decrease the severity, but not the incidence of stunting. Household targeting could reduce the malnutrition gap by over 20 percent and the quadratic malnutrition gap by over 30

percent, but could be very expensive to implement. 'Broad stroke' geographic targeting could reduce the same measures by 15 percent and 20 percent respectively, and would be cheaper to implement.

- Education poverty: In the field of education, one can for instance use illiteracy among children in age of primary or secondary school as the characteristic identifying the education poor. Alternatively, in countries where literacy is high among new cohorts of children, one can compare the actual number of years of education completed by each child to a "poverty line" equal to the expected number of years of education for a child of that age without repetition and drop-out. The results of such an exercise for Panama suggest that 29 percent of all children in age of primary school in urban areas are behind in their level of schooling in terms of age-for-grade, versus 51 percent in rural areas. This figure is the equivalent of the headcount index for traditional poverty measurement. The "headcount" of education poverty among children in age of secondary school is much higher, at 60 percent for urban areas and 89 percent in rural areas, because delays in completing primary education carry on to the secondary cycle and some students do not pursue their studies beyond the primary cycle. When the number of years of education separating the children from what they should have achieved at their age under perfect schooling conditions is taken into account, one can compute poverty gaps and squared poverty gaps.

- Energy or fuel poverty: Many households in developing countries cannot satisfy their basic energy needs. Although fuel is not a dimension of well-being as such, it is an important input for a healthy life (through cooking and heating), for a better education (lighting), and for higher security. One could define 'fuel poverty' as the inability by households to meet basic energy needs (in Guatemala, the energy required to run two 60 watt light bulbs and a 16 watt radio for four hours and ten kilograms of fuel wood each day were taken as the basic needs – equivalent to 5.9 Kilowatt-hour every day). Table 28 shows that households with access to electricity consume on average 3,804 Kwh per year and pay 0.52 Qz per kilowatt-hour on average. Meanwhile, households without access to electricity have a lower energy consumption (2.892 kwh) and pay a higher price (1.35 Qz/kwh). One fourth of the population with access to electricity is fuel poor (headcount of 25.5%), as compared to half of the population without access (headcount of 0.509). Using estimates of the impact of access to electricity on the average price paid for energy, the authors estimate that if households without access were given access to electricity, the headcount among that group would be reduced to 36.5%.

Table 28: Fuel poverty with and without access to electricity in Guatemala, 1998/99

	Households with access to electricity	Households without access to electricity	
		Current situation	With access
Net consumption (all energy sources, kwh)	3804	2892	3967
Price per effective kwh	0.52	1.35	0.98
Fuel poverty headcount (%)	25.5	50.9	36.5

Source: Foster, Tre, and Wodon (2000).

TN 7 Inequality measures and their decompositions

Inequality measures have been introduced in the main text. This note provides mathematical expressions for the three main measures: the Gini, Theil, and Atkinson indices. Each index can be generalized in order to put more weight on selected parts of the distribution of income or

consumption. As is the case for poverty measures, some inequality measures can be decomposed, and this note presents decomposition formulas for the GE class which includes the Theil index.

Inequality measures

The standard Gini index measures twice the surface between the Lorenz curve, which maps the cumulative income share on the vertical axis against the distribution of the population on the vertical axis, and the line of equal distribution (see Figure 4 in main text, **section 3.1**). A large number of mathematical expressions have been proposed for the Gini index, but the easiest to manipulate is based on the covariance between the income Y of an individual or household and the F rank that the individual or household occupies in the distribution of income (this rank takes a value between zero for the poorest and one for the richest). Denoting by $\bar{y}$ the mean income, the standard Gini index is defined as:

$$Gini = 2\ cov\ (Y, F) / \bar{y}$$

The Gini has attractive theoretical and statistical properties which other inequality measures do not have, which explains why it is used by most researchers. For a review and discussion of these properties, see the chapter on **Inequality and Social Welfare**. The extended Gini uses a parameter ν to emphasize various parts of the distribution. The higher the weight, the more emphasis is placed on the bottom part of the distribution (ν=2 for the standard Gini index):

$$Gini(\nu) = \frac{-\nu \operatorname{cov}(y, [1-F]^{\nu-1})}{\bar{y}}$$

Another family of inequality measures is the **General Entropy measure**, defined as:

$$GE(\alpha) = \frac{1}{\alpha^2 - \alpha}\left[\frac{1}{n}\sum_{i=1}^{n}\left(\frac{y_i}{\bar{y}}\right)^{\alpha} - 1\right]$$

With $GE(0) = \frac{1}{n}\sum_{i=1}^{n}\log\frac{\bar{y}}{y_i}$, $GE(1) = \frac{1}{n}\sum_{i=1}^{n}\frac{y_i}{\bar{y}}\log\frac{y_i}{\bar{y}}$ and $GE(2) = \frac{1}{2n\bar{y}^2}\sum_{i=1}^{n}(y_i - \bar{y})^2$

Measures from the GE class are sensitive to changes at the lower end of the distribution for α close to zero, equally sensitive to changes across the distribution for α equal to one (which is the Theil index), and sensitive to changes at the higher end of the distribution for higher values.

A third class of inequality measures was proposed by Atkinson. This class also has a weighting parameter ε (which measures aversion to inequality) and some of its theoretical properties are similar to those of the extended Gini index. The Atkinson class is defined as follows:

$$A_{\varepsilon} = 1 - \left[\frac{1}{n}\sum_{i=1}^{n}\left(\frac{y_i}{\bar{y}}\right)^{1-\varepsilon}\right]^{1/(1-\varepsilon)}$$

Decomposition of inequality measures: Illustrations for the GE class

Inequality is often decomposed by population groups to assess the contribution to total inequality of inequality within and between groups, for instance within and between individuals in urban and rural areas. Inequality measures can also be decomposed according to consumption or income sources in order to identify which component contributes most to overall inequality. Finally, decompositions can be used to analyze changes in income inequality over time. Below, decompositions are provided for the GE class. The chapter on **Inequality and Social Welfare** includes a detailed discussion of applications of decompositions of the extended Gini class of inequality measures which have especially attractive properties for policy simulations.

Decompositions at one point in time

Total inequality *I* can be decomposed into a component of inequality between the population groups I_b and the remaining within-group inequality I_w. The decomposition by population subgroups of the GE class is defined as:

$$I = I_w + I_b = \sum_{j=1}^{k} v_j^{\alpha} f_j^{1-\alpha} GE(\alpha)_j + \frac{1}{\alpha^2 - \alpha}\left[\sum_{j=1}^{k} f_j\left(\frac{y_n}{\bar{y}}\right)^{\alpha} - 1\right]$$

where f_j is the population share of group *j* (*j=1,2,..k*); v_j is the income share of group *j* ; and y_j is the average income in groups *j.*

Inequality measures can also be decomposed by source of consumption or income. The decomposition for the GE measure with α=2 is as follows:

$$I = \sum_f S_f = \sum_f \rho_f \frac{\mu_f}{\mu} \sqrt{GE(2).GE(2)_f}$$

where S_f is the contribution of income source *f;* ρ_f is the correlation between component *f* and total income; and μ_f / μ is the share of component *f* in total income. If *Sf* is large, then component *f* is an important source of inequality.

Decompositions for changes in inequality over time

Using sub-group decompositions, changes in inequality can be decomposed into: 1) changes in the numbers of people in various groups or "allocation" effects; 2) changes in the relative incomes of various groups or "income" effects; and 3) changes in inequality within groups or "pure inequality" effects. Because the arithmetic can be complex for some inequality measures, this decomposition is usually applied only to Generalized Entropy index GE(0) as follows:

$$\Delta GE(0) = \underbrace{\sum_{j=1}^{k} \bar{f}_j \Delta GE(0)_j}_{\text{Pure inequality effects}} + \underbrace{\sum_{j=1}^{k} \overline{GE(0)_j}\, \Delta f}_{\text{Allocation effects}} + \underbrace{\sum_{j=1}^{k} \left[\bar{\lambda}_j - \overline{\log(\lambda_j)}\right] \Delta f_i}_{\text{Allocation effects}} + \underbrace{\sum_{j=1}^{k} \left(\bar{v}_j - \bar{f}_i\right) \Delta \log(\mu(y))_j}_{\text{Income effects}}$$

where Δ is the difference operator, λj is the mean income of group j relative to the overall mean (i.e., λj = μ(yj)/μ(y)) and the over-bar represents averages. The first term captures the pure inequality effects, the second and third terms, the allocation effects, and the fourth term, the income effects.

Using source decompositions, changes can be decomposed by income source. This allows to see whether an income source *f* has a large influence on changes in total inequality over time. For the General Entropy index with α=2, defining S_f as above, the decomposition is:

$$\Delta GE(2) = \sum_f \Delta S_f$$

TN 8 Using linear regressions for analyzing the determinants of poverty

It has become a standard practice to analyze the determinants of poverty through categorical regressions such as probits and logits (see Box 7 in the main text). When using such categorical regressions, it is assumed that the actual (per capita) income or consumption of households is not observed. We act as if we only know whether a household is poor or not, which is denoted by a categorical variable which takes the value one if the household is poor, and zero if the household is not poor. Under the hypothesis of a normal standard distribution for the error term, the model is estimated as a probit. If the error term is assumed to have a logistic distribution, the model is estimated as a logit. The main problem with categorical regressions is that the estimates are sensitive to specification errors. With probits, the parameters will be

biased if the underlying distribution is not normal. More generally, the model does not make use of all the information available, because it collapses income or expenditure into a binary variable. This does not mean that probit or logit regressions should never be used. Categorical regressions will typically have better predictive power for targeting, that is for classifying households as poor or non-poor (see **Technical Note 9**).

The alternative is to use the full information available for the dependant variable (indicator of well-being), and to run a regression of the log on the indicator (if the distribution is log normal.) Assume that w_i is the normalized indicator divided by the poverty line, so that $w_i = y_i/z$, where z is the poverty line and y_i is (per capita) income or consumption. A unitary value for w_i signifies that the household has its level of income or consumption exactly at the level of the poverty line. Denoting by X_i the vector of independent variables, the following regression can be estimated:

$$\text{Log } w_i = \gamma' X_i + \varepsilon_i$$

From this regression, the probability of being poor can then be estimated as follows:

$$\text{Prob}[\log w_i < 0 \mid X_i] = F[-(\gamma' X_i)/\sigma]$$

where σ is the standard deviation of the error terms and F is the cumulative density of the standard normal distribution.

Once regressions have been estimated to analyze the determinants of poverty, the coefficients on the variables (γ) can inform on the various correlates of poverty and be used to simulate the impact of various policies. Figure 9 presents an Excel Dialog Box which was prepared for Bolivia using the coefficients of multivariate regressions on the determinants of the logarithm of per capita income. Apart from a constant, the independent variables in the regressions (which are done separately for urban and rural areas) include (a) the geographic location of the household according to Bolivia's main cities and departments in rural areas; (b) household level variables, including the number of babies, children, and adults and their square, whether the household head is a woman, the age of the head and its square, the marital status of the head, the migration status of the head (since birth and/or in last five years), and whether the household head speaks one of the main indigenous languages; (c) characteristics of the household head, including his/her level of education; whether he/she is unemployed and searching for work, not working, and has a secondary occupation apart from his/her primary occupation; his/her sector of activity; his/her position; whether he/she works in the public and/or formal sector; the size of the firm in which he/she works; and whether he/she has been sick and for how long; and (d) the same set of characteristics for the spouse of the household head, when there is one. The Dialog Box can be used to run simple simulations of the impact of any one of these variables on per capita income and the probability of being poor or extremely poor, holding constant all other variables. For example, one can see what is the impact of raising the education level of the head of a given household has on income and poverty. The user can also test how the probability of being poor or extremely poor changes with the choice of the poverty line.

Figure 9: Using regression estimates to build easy-to-use excel software for simulations

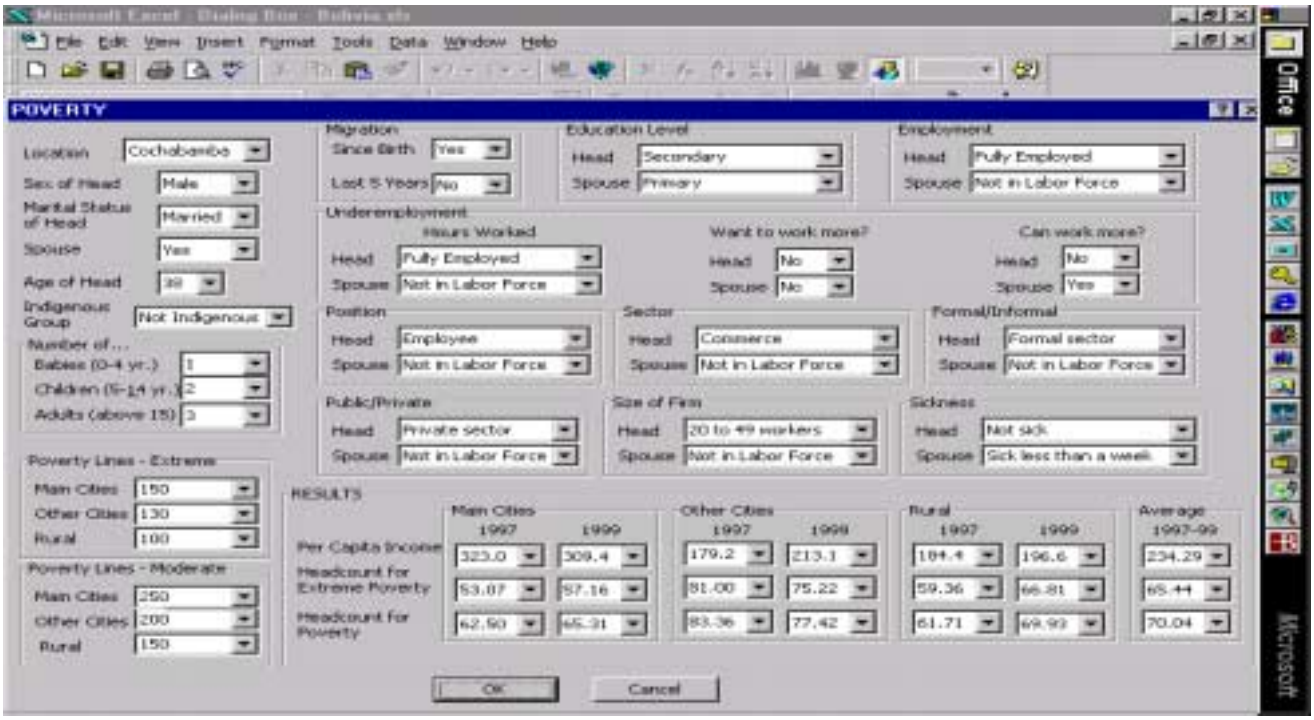

Source: Wodon (2001a)

TN 9 Using categorical regressions for testing the performance of targeting indicators

While categorical regressions should in most cases not be used to analyze the determinants of poverty, they can be used to measure the performance of alternative targeting indicators for social programs and transfers. This is an important issue, because Governments often use proxies for assessing the level of income or consumption of a household or individual in order to assess eligibility for programs benefits. Categorical regressions can help in choosing the best proxies for identifying the poor and the non-poor, or more broadly for selecting beneficiaries. ROC (Receiving Operating Characteristics) analysis is one useful technique based on an underlying logit regressions of the probability of being poor. It can be implemented easily with statistical packages such as STATA. The benefits of ROC analysis for selecting targeting indicators or proxies for income or consumption are discussed by Wodon (1997b).

To explain the principles behind ROC analysis, denote by P, P^-, and P^+ the number of the poor, the number of the poor classified as non-poor and the number of the poor classified as poor by a given econometric model. Also denote by NP, NP^-, and NP^+, the number of the non-poor, the number of the non-poor classified as non-poor and the number of the non-poor classified as poor. Sensitivity $SE = P^+/(P^- + P^+) = P^+/P$ is the fraction of poor households classified as poor. Specificity $SP = NP^-/(NP^- + NP^+) = NP^-/NP$ is the fraction of non-poor households classified as non-poor. The probability of errors of type I (identifying as poor a non-poor household or individual) and type II (identifying as non-poor a poor household or individual), as they are usually defined in economics, are captured by $(1- SP)$, and $(1-SE)$.

Table 29: Terminology used in ROC analysis for targeting purposes

Predicted Status	Actual Status: Non-Poor	Actual Status: Poor
Non-Poor	$SP=NP^{-}/(NP^{-} + NP^{+})$	$1 - SE= P^{-}/(P^{-} + P^{+})$
Poor	$1 - SP=NP^{+}/(NP^{-} + NP^{+})$	$SE=P^{+}/(P^{-} + P^{+})$

When using a statistical package and running a probit or logit regression for poverty, each observation is given an index value equal to the predicted right hand side of the regression. This predicted value is used to classify the households as poor or non-poor, with the computer typically using one half as the cut-off point (those above the cut-off point are classified as poor). But this cut-off point can be changed. A ROC curve is a graph that plots *SE* as a function of 1 – *SP* for alternative values of the cut-off point. Figure 10 shows an hypothetical ROC curves. The higher the ROC curve, the better the predictive power of the model used for making the predictions. A 45 degree line has no predictive power while a vertical line from the origin to the top of the box followed by a horizontal line until the upper right corner has perfect predictive power. The area below a ROC curves provides a summary statistic of the predictive value of the underlying model used for targeting. An area of 0.5 corresponds to the 45 degree line which has no explanatory power. An area of one corresponds to perfect prediction.

Figure 10: An hypothetical ROC curve

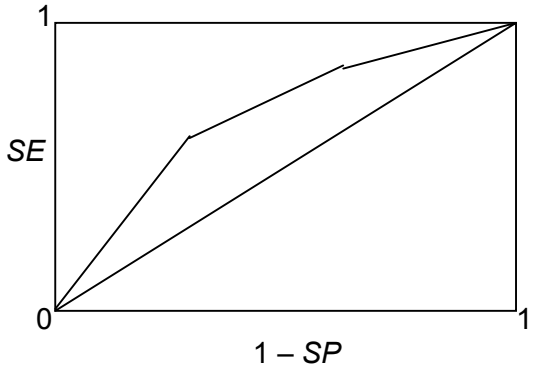

If the ROC curve of one targeting indicator (or set of indicators) used to predict poverty lies above the ROC curves of all the alternatives, that indicator will typically be the best to target the poor for the class of social welfare functions based on the two types of errors that can be committed through targeting. If two ROC curves intersect, the choice of the best indicator will depend on the normative weights attached by the policy maker to the two types of errors.

ROC analysis was used for example by Estache et al. (2001) to assess how well various indicators performed for identifying the poor in order to target electricity subsidies among households with a connection to the public electricity grid in Honduras. The results are given in Table 30 which provides the areas under the ROC curves for the various targeting indicators. The best results (i.e., the largest areas under the ROC curve) are obtained using a combination

of different characteristics. As far as single characteristics go, electricity consumption has some predictive power, but less so than some other variables, such as the size or quality of the house.

Table 30: Areas under ROC curves for alternative targeting indicators in Honduras

	Performance in Identifying the extreme poor	Performance in identifying the poor
Socioeconomic status (multiple characteristics)	87%	83%
Demographics	72%	71%
Educational attainment	71%	72%
Employment status	69%	66%
Geographic location (department)	66%	63%
Housing characteristics (multiple characteristics)	82%	81%
Size of house	77%	77%
Quality of house	72%	72%
Access to electricity	68%	69%
Access to water and sanitation	61%	58%
Electricity consumption	70%	73%

Source: Estache et al. (2001).

Of course, the choice of targeting mechanisms for programs and subsidies depends on several other factors, including administrative targeting costs and political economy considerations. A more detailed a discussion of targeting issues is provided in the chapter on **Public spending**.

TN 10 Using wage and labor force participation regressions

Similar to the analysis of correlates of poverty, regressions can be used to analyze the determinants of individual labor income. To analyze the impact of individual characteristics on labor income, and to measure among other things the impact of a better education on earnings, other types of regressions must be used. The standard approach consists in running a so-called Heckman model. Denote by log w_i the logarithm of the wage (or earnings) observed for individual i in the sample. The wage w_i is non zero only if it is larger than the individual's reservation wage (otherwise, the individual chooses not to work.) The difference between the individual's wage and reservation wage is denoted by Δ^*_i. The individual's wage on the market is determined by geographic location (separate regressions are run for the urban and rural sectors), years of experience E, and years of schooling S. There may be other determinants of wages but these are not observed. The difference between the individual's wage and his reservation wage is determined by the same characteristics, plus the number of babies B, children C, and adult family members A of the individual (and their square.) The Heckman model is written as:

$$w_i = w^*_i \text{ if } \Delta^*_i > 0, \text{ and } 0 \text{ if } \Delta^*_i < 0$$

$$\text{Log } w^*_i = \alpha_w + \beta_{w1}E_i + \beta_{w2}E_i^2 + \beta_{w3}S_i + \beta_{w4}S_i^2 + \varepsilon_{wi}$$

$$\Delta^*_i=\alpha_\Delta+\beta_{\Delta1}E_i+\beta_{\Delta2}E_i^2+\beta_{\Delta3}S_i+\beta_{\Delta4}S_i^2+\beta_{\Delta5}B_i+\beta_{\Delta6}B_i^2+\beta_{\Delta7}C_i+\beta_{\Delta8}C_i^2+\beta_{\Delta9}A_i+\beta_{\Delta10}A_i^2+\varepsilon_{\Delta i} = m_{\Delta i} + \varepsilon_{\Delta i}$$

The expected value of ε_{wi} is not zero. Denoting by φ and Φ the standard normal density and cumulative density, and noting that σ_Δ, the standard error of $\varepsilon_{\Delta i}$, is normalized to one, we have:

$$E[\text{Log } w^*_i \,|\Delta^*_i>0] =\alpha_w+\beta_{w1}E_i+\beta_{w2}E_i^2+\beta_{w3}S_i+\beta_{w4}S_i^2+\lambda\varphi(m_{\Delta i})/\Phi(m_{\Delta i})$$

$$E[\text{Log } w^*_i \,|\Delta^*_i<0] =\alpha_w+\beta_{w1}E_i+\beta_{w2}E_i^2+\beta_{w3}S_i+\beta_{w4}S_i^2-\lambda\varphi(m_{\Delta i})/[1-\Phi(m_{\Delta i})]$$

If λ is statistically different from zero, the returns to education will differ between the employed and the unemployed, although the difference will typically be small. Simple approximations of the private returns to education (or more precisely, of the marginal impact of a better education on individual earnings) can be computed from the above wage regressions by taking the first derivative of the expected wage with respect to the number of years of schooling. Thus the "return" to education for year of schooling S is $\partial E[\text{Log } w^*_i]/\partial S = \beta_{w3}+2\beta_{w4}S$ when λ is zero. The returns are increasing (decreasing) with the number of years of schooling if the coefficient β_{w4} is positive (negative.) These returns do not take into account the positive impact on the probability of working of education (i.e., the fact that $\beta_{\Delta 3}S_i+\beta_{\Delta 4}S_i^2$ is typically positive.) The returns also do not include estimates of the costs of schooling for parents and society (which reduce the returns) and of the indirect effects and externalities associated with education (which typically increase the returns, from the point of view of both the society and the household.) For more information on these techniques, see the chapter on **Education**.

TN 11 Limitations of income variability analysis

As explained in the main text, income variability analysis based on standard deviations (SDs) and coefficients of variation can provide some information on the vulnerability of households, yet it has limitations. These can be illustrated with a simulation of various types of households. In Figure 11 below, each group has the same **average** income over the ten periods (100), but the patterns of their income over time differ greatly. Clearly, Groups B, C, and D are more vulnerable than Group A. This is reflected in their higher SD.

Figure 11: **Comparison of Household Vulnerability**

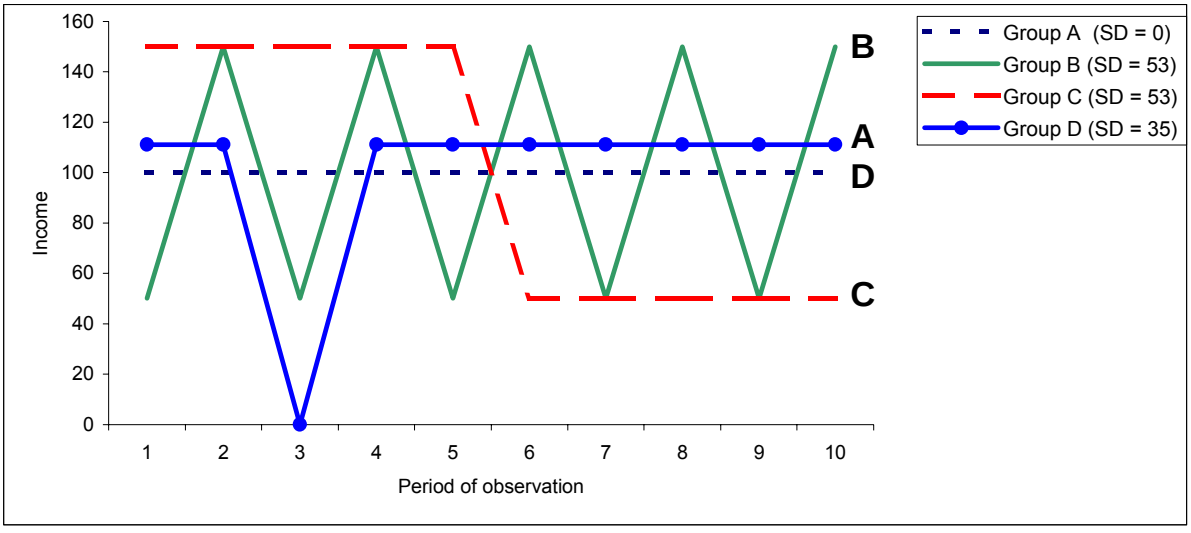

Comparing Group B and Group C shows that their SD is similar but their patterns are different. Imagine that B is a rural household with the fluctuations following peak and mean periods. C is an urban household that experiences the death of its main breadwinner in period 5, resulting in a permanent loss of income. Its vulnerability is different but this is **not** reflected in the SD.

Turning to Group D, imagine that the wage earner of the household has experienced an acute but temporary spell of illness, or a short time of unemployment, in period 3. If the household had stocks or buffers, it probably could have used these to maintain its consumption, but in the opposite case, it could have suffered irremediable losses during that one period, such as the death of a child from under-nutrition or the stalled mental development of children. The SD indicates a lower vulnerability for this household than for that of Group B, which is not necessarily true.

TN 12 Beyond poverty: Extreme poverty and social exclusion

There is wide agreement that poverty is multidimensional. In order to go beyond the emphasis placed on income and consumption in analytical work on poverty, this note briefly introduces the concepts of extreme poverty and social exclusion, and how these concepts differ from monetary (i.e., income or consumption-based) poverty as it has been traditionally defined.

Extreme poverty

While it is certainly feasible to define extreme poverty as a very low level of income or consumption, this note suggests another interpretation based on the grass-roots work of NGOs. According to Wresinski (1987; see also Wodon 2001b), three distinctions can be made between poverty and extreme poverty. First, extreme poverty results from a lack of various "basic securities" or assets in many areas of life (education, health, employment, etc.), rather than in the income or consumption space only. This lack of several basic securities may have a cumulative impact and lead to an insecurity affecting new dimensions in a poor person's life. This can be analyzed through interaction effects. Second, extreme poverty tends to be associated with the persistence of this insecurity over possibly long periods of time. Third, the extreme poor are often unable to exercise their rights and assume their responsibilities, which has operational implications when conditions are imposed for participation in social programs.

- Interaction effects: Various factors or dimensions of poverty may reinforce each other (interaction effects) and prevent the very poor from emerging from extreme poverty. In other words, extreme poverty results not only from a lack of financial resources, but also from a lack of education, employment, housing, health care, as well as civil and political rights. Beyond some threshold, the lack of many different basic securities has mutually reinforcing impacts. The poor become extremely poor and are prisoners of a vicious circle.

- Long term dimension: The time dimention associated with poverty has been discussed in the section devoted to vulnerability. Extreme poverty often persists through time, and tends to be transmitted from one generation to the next. That is, a common feature among the extreme poor is the permanence, or at least the recurrence of their situation. One can think of extreme poverty as being persistent, or at least chronic. The longer the experience of poverty, the harder it is to emerge out of extreme poverty. This is difficult to analyze through panel data, but can be revealed through qualitative work, such as life histories (see **Technical Note 13**).

- Rights and responsibilities: The third reference in the Wresinski approach to extreme poverty deals with rights and responsibilities, and rests on two articulations. The first articulation highlights, for each human right at a time, the existing link between the access to that right and the exercise of a corresponding responsibility. This link is broken when, due to a lack of access to the right, the individuals or families in poverty cannot fulfill their

corresponding responsibility. In turn, because they cannot demonstrate their ability to fulfill their responsibility, the poor are not in a position to claim their right. For example, a person who has been unemployed for a long period of time will have difficulties to acquire the credentials that would enable him/her to demonstrate his/her ability to work. The second articulation refers to the interdependence or indivisibility between various rights. For example, without access to one right, it is difficult to exercise other rights.

Rights and social exclusion

In a series of recent resolutions, various agencies of the United Nations system have suggested that extreme poverty may lead to violations of human rights in their indivisibility. Still, the use of the discourse of human rights to analyze poverty is not as well accepted as the quantitative tools described in the main text of this chapter. The reference to human rights is not meant to be dogmatic. Rather, it is meant to provide a framework for the discussion of some of the institutional issues related to poverty policies. In Freeden's (1991) terms, a human right is a conceptual device that helps societies to assign a priority to certain human or social attributes regarded as essential to the adequate functioning of a human being. In other words, from a conceptual point of view, at least two requirements are needed for a basic security or need to be elevated to the status of human right. Human rights must be recognized as such by others than their most direct beneficiaries, and they must be essential to our functioning as human beings. These two requirements are linked to each other since only key human attributes stand a chance to benefit from a consensus for their recognition as human rights by society.

Only society or the community as a whole can grant a right to someone. Conversely, the non-attribution of a right expresses a sanction, an exclusion, whether it is explicit or implicit, of some individuals. In a report on social exclusion and anti-poverty strategies, the International Institute for Labor Studies (1997) argues that the concept of social exclusion provides an integrated and dynamic analytical perspective with which to analyze the relationships between well-being and rights. When the non-attribution of a right is explicit, the exclusion from the benefit of the protection granted by a right is operated through institutional procedures. If the rules for the enjoyment and the exercise of human rights skim some beneficiaries, thereby ending in denials of rights for those who cannot comply with society's rules, institutional change may be needed in order to fight extreme poverty.

Rather than opposing the traditional concepts poverty with those of extreme poverty and social exclusion, and arguing that one set of concepts is more useful than the other, it is better to consider extreme poverty and social exclusion as complementary concepts to that of poverty. Material (consumption or income) poverty is then viewed as a particular form of deprivation, i.e. the lack of command over goods and services. Social exclusion is broader than poverty in considering issues of social participation and rights realization, as well as processes. Extreme poverty is deeper than poverty in that it results from many handicaps faced by the very poor.

Still, one important distinction between the concepts of poverty or extreme poverty, and that of exclusion, is that poverty and extreme poverty are states of well-being, while exclusion is a dynamic process whereby someone excludes someone else. Literally, to exclude means to banish, to send somebody away from a place where she had the right to stay before. By extension, to exclude means to deprive somebody from any right previously granted, or normally granted, to those recognized as full members of a given community or society. The strength of an analysis based on a social exclusion perspective is to put the issues together into a cohesive framework, that is to allow the analyst to investigate the causes and processes involved in the persistence of poverty, often using both quantitative and qualitative techniques.

A social exclusion perspective can thus function as a flexible analytical framework aimed at understanding social disadvantage (Gacitua-Mario, Sojo, and Davis 2000; Gacitua-Mario and Wodon 2001). Beyond being "goods-centered", the analysis of social exclusion is "people-centered" and "institutions-centered". The analysis of social exclusion also contains both an objective and subjective dimension, to the degree that it considers both the objective conditions of people's lives and their perceptions of being connected or disconnected from wider spheres of social, political and cultural life. This perspective does not substitute for the traditional income- or consumption-based notion of poverty or other concepts dealing with vulnerability. It rather provides a framework for interrelating different levels of analysis (multidimensionality) and cumulative processes that maintain or pull social groups into social disadvantage. The perspective fully recognizes the importance of the traditional dimensions of poverty, such as the inability to generate a sufficient and stable income and to have access to quality social services in order to meet basic needs. However, it also incorporates other dimensions that belong to the relational/symbolic domain: the socio-organizational cultural and the political. The value added by a social exclusion perspective lies in the emphasis on dynamic processes where both institutions and agents are involved. For institutions, this refers to the ways institutions, rules and perceptions interact to generate or combat processes of social disadvantage (see also the chapter on **Community Development**).

TN 13 Qualitative and Participatory Assessments

The discussion in the preceding note naturally leads to qualitative and participatory research methods. Among these, so-called Participatory Poverty Assessments (PPAs) are tools for consulting the poor directly and systematically. Using PPAs can deepen the understanding of poverty, explain processes of impoverishment and household survey data, convey the priorities of the poor, and assist in analyzing poverty beyond the household unit. PPAs can capture dimensions of poverty that are not always addressed in household surveys. Experience has shown that poor people speak of poverty in different terms than those typically used in policy analysis. They may refer to such characteristics as vulnerability, physical and social isolation, lack of security and self-respect, powerlessness, and lack of dignity. In addition, past PPAs have offered insights into dimensions of poverty that had not been previously examined by household surveys, such as vulnerability, gender, crime and violence, and seasonality.

The types of qualitative data that are important for poverty reduction strategy design and monitoring include the following:

- Poor people's priorities for improving their situation—desegregated by sex and by other important characteristics of the community
- Causality data—people's perceptions of causes and consequences of poverty
- Opportunities poor people see for improving their situation
- Constraints and barriers to improving their situation
- Locality differences—differences between districts and between rural and urban centers
- Perceptions on quality of service delivery, infrastructure and governance at the local level
- Identification of who the poor are

PPAs use a variety of methods that combine visual techniques, such as mapping, matrices, and diagrams, with verbal techniques, such as open-ended interviews and discussion groups. The process of undertaking a PPA is different at the community and national levels. The community level involves undertaking the study with the locally led teams. The national-level study involves linking communities to a broader policy dialogue.

PPAs at the Community Level

Many PPAs use the participatory rural appraisal (PRA). The PRA offers tools such as mapping; diagrams of changes, trends, and linkages; matrices; and scoring. It is a locally led effort and requires a willingness to unlearn assumptions and conditioned responses. The PPA design will depend on the country context, research agenda, sample, and researchers' experience. Table 31 describes characteristics of the average PPA.

Table 31. Design Features for Participatory Poverty Assessments (PPAs)

Feature	Detail
Cost	$75 – 125,000
Number of communities selected for research	40 – 60 communities
Time spent on training	2 weeks
Time spent on field research	3 – 6 months
Time spent on analysis	2 – 3 months
Size research team (including team leaders and trainers)	10 – 20 people
Composition of research team	Country nationals, half men and half women, ability to speak local languages, representatives from various ethnic groups and a cross-section of age groups.
Typical agency conducting the field work	Government extension workers; local and international NGOs; academic institutions; independent consultants and firms
Donors who have contributed to government led PPAs	DFID, World Bank, Action Aid, Oxfam, UNDP, UNICEF, DANIDA, Asian Development Bank

PPAs and household surveys can inform each other, so the sequencing will be determined by the context in- country. If the PPA comes first, its results can help focus the research agenda for the quantitative survey and generate hypotheses. Conversely, the results of quantitative surveys can be used to identify the poorest geographical areas on which participatory research should focus or can be used to identify specific set of issues that require further understanding. Good practices often have iterative processes.

The Zambian PPA (see case examples below) illustrates how issues of measuring poverty, and analyzing the behavior of and impact of policies on the poor, can be addressed using qualitative techniques. In particular, issues in rural and urban areas and differences in attitude and behavior between men and women can be explored.

The Process at the National Level

Participatory policymaking links communities to a broader policy dialogue that includes a cross-section of stakeholders. In general, open political environments provide greater opportunities for building consensus on poverty issues. In Costa Rica, for example, where there is a tradition of bringing marginal groups into the political sphere, the government was eager to better understand poverty from the perspective of the poor and welcomed the PPA. See case examples (below) for illustrations.

Experience has shown that involving key policymakers from the beginning enhances ownership and commitment. Limited government support, or lack of government support, can hinder the impact of the PPA, especially if the research results run counter to the government's interest. Generating a more open climate and starting with small-scale participation and PPAs initially

can help make the process of participation less daunting. Where appropriate, the following measures can help increase policy impact:

- Involve policymakers in the early planning of the PPA.
- Start with small-scale participation in constrained political environments.
- Bring key policymakers to the field to participate in the PPA research.
- After the results are presented, convene workshops with policy-makers and local people.
- Negotiate high-level commitment to follow up the PPA and monitor the implementation of key recommendations.

Case Examples

PPA Highlights the Potential of Women's Groups in Kenya

The coping strategies of the poor, most of whom do not have access to credit, depend on diversifying their livelihoods and on the strength of their social networks and informal groups. Because their livelihoods are so diversified, no single employment program will reach the poor. The informal groups and associations, on the other hand, engage in a wide range of economic and social welfare activities. The PPA in Kenya highlighted the untapped potential of these groups to reach the poorer segments of society. The study estimated that at least 300,000 groups and associations exist in rural Kenya, including more than 23,000 registered women's groups. Every village was found to have from five to seventeen different types of groups, and more than one active or defunct women's group. The following are some of the findings that emerged about these women's groups:

- During discussions of coping strategies, women's self-help groups were mentioned frequently in every district, in particular by female-headed households.
- In addition to income generation, group objectives frequently included welfare activities: raising cash to pay school fees, meet hospital expenses, or help with transport costs to bring the dead back to the villages for burial.
- Most groups levied membership fees and monthly contributions.
- Although high fees excluded the poor, many groups targeted their activities specifically to assist the poor with food, school fees, and housing construction.
- Women's groups were often formed along clan or kinship lines and often had male members. Generally, they were supported by village men and the community at large.

Based on the findings of the PPA, proposals to reach the poor by strengthening women's groups include legal registration so that groups are eligible for credit, technical and business management training of group members, and extension of micro-enterprise credit to groups.

Zambia's Participatory Poverty Assessment: Objectives and Method

The key goals of the PPA were to explore local concepts of poverty, vulnerability, and relative well-being in poor urban and rural communities, what the poor themselves see as the main concerns and problems and effective actions for poverty reduction, and investigate local perceptions of key policy changes for economic reform. The PPA included a mix of qualitative techniques, such as unstructured and semi-structured interviews, focus group interviews, wealth and well-being ranking, institutional diagrams (Venn diagrams), and seasonality diagramming. Ten research sites were selected to represent a variety of communities, including rural, urban, cultural or ethnic group, and so on. The issues and techniques are presented in Table 32.

Table 32. Issues Addressed and Qualitative Techniques Employed in the Zambian PPA

Poverty Issue	Qualitative methods used
Perceptions and indicators of wealth, well-being, poverty, vulnerability, powerlessness. Local terminology and correspondence with such concepts. Differences in perception by gender.	• Wealth and well-being ranking or grouping for criteria and indicators • Semi-structured interviews • Social mapping
Perception of change over time in welfare, indicators, terms of trade	• Timelines (for migration, terms of trade, environment, and so on)
Access to (and use of) services such as health, education, credit; preferences especially where choice between options is available; perceptions of services, including views (or awareness) of recent change; differing perceptions and values for men and women	• Institutional diagramming • Semi-structured interviews • Trend analysis of services, for example, health, education, agricultural extension, marketing
Seasonal stress: food security, health, general livelihoods; income, spending, activity (by selected occupational groups)	• Seasonal calendar (health, food security, food intake, access to fuel, water, and so on) • Comparative seasonal calendars (good years, bad years, average years)
Assets of rural communities--access to services, common property resources, other natural resources	• Resource mapping • Focus groups • Institutional (Venn) diagramming
Assets of households	• Wealth ranking or grouping • Social mapping • Semi-structured interviews
Coping strategies and fallback strategies in times of crisis	• Livelihood analysis • Semi-structured interviews • Ranking exercises
Perception of consumption level in terms of food, clothing, and relation to well-being	• Well-being grouping/ranking • Social mapping • Semi-structured interviews
Community-based support mechanisms for the rural poor (community safety nets)	•
Local institutions of self-help and support for the urban poor (for example, market traders' association, trade associations, churches, and so on)	• Semi-structured interviews • Institutional mapping
Role of community institutions in service and infrastructure provision	• Institutional mapping • Semi-structured interviews
Long-term environmental trends, for example., declining soil fertility, declining rainfall	• Historical transects • Community timelines • Resource mapping at different points in time • Trend analysis
Responsibilities, obligations within households (support to children, provision of food, payment of school fees, and so on, by gender)	• Semi-structured interviews • Decision-making matrix

Source: World Bank (1994)

Assessing existing qualitative data

Desk Review. A desk review would be warranted if qualitative data sources cannot be readily identified and the data are not routinely used for poverty profiling and poverty reduction strategies. A desk review would not be needed if it had already been performed in the past few years. The review would identify and summarize qualitative data from official poverty assessments, NGO reports, participatory research, and needs assessment processes. Such data would be assessed for timeliness, quality, coverage and depth, that is, how well the data help policymakers to derive policy. Table 33 below shows criteria for judging the quality of existing qualitative data:

Table 33: Criteria for Assessing Adequacy of Qualitative Data

Criteria	Adequate requirement
1. Age of data	• Collected in the past five years
2. Methodologies	• Participatory methods (PRA), focus groups better
3. Coverage and scope: • Geographical • Rural and Urban • Groups consulted	• All major agro-ecological zones represented • Both • Both sexes, youth and elderly, other vulnerable groups, major livelihood groups of the poor
4. Dimensions of poverty	• Dynamics (esp. seasonality), causality, gender, age, livelihood • Identification of vulnerable groups
5. Perceptions of services, infrastructure and governance	• Explored
6. Information flows	• Awareness and understanding of poverty-related policies and programs
7. Priorities of the poor	• Opportunities and constraints improving quality of life, priorities for poverty reduction

TN 14 Use of Demographic and Health Surveys for Poverty Analysis

While demographic and health surveys have not been used very much for the analysis of well-being, they go far beyond collecting health information. For example, the surveys cover:

(a) *Basic service access*: source of water, how far to get to water, electricity access, type of toilet facilities, materials used for the floor, ownership of durable goods (for example, car, motorcycle, bicycle, radio).

(b) *Education*: highest formal education attained of all household members; reason why women stopped attending school; current school enrollment, but not by school type, of all household members.

(c) *Occupation of adults*.

(d) *Migration*: residence of household.

(e) *Health*: infant mortality, fertility, contraception practices and family planning, health attendance during pregnancy, feeding practices, vaccination of children, child illnesses (below 5 years of age), knowledge about disease treatment, sickness and health center use of mothers and children, satisfaction with health service, cost of treatment; knowledge about age; female circumcision; height and weight of children.

The DHS surveys do not contain household income or consumption but "wealth" information can be used to allow derivation of a poverty profile. A household wealth indicator can constructed using available information on durable goods, basic services, and so on. The indicator is then ranked to construct quintile distributions. There are currently several different methods to derive such a wealth indicator, but the profiles tables constructed using the indicator can include:

(a) Distribution of poverty (or rather, low wealth), basic service access, education and health by region, since the survey is able to produce estimates for many regions.

(b) Profile of health outcomes, access to basic services, health and education by wealth quintile (this can also be done by sector – urban versus rural – and by gender).

This information gives a good base to see whether pockets of (wealth) poverty correlate with health and education outcome indicators, whether the infrastructure and educational deficit in these areas is particularly high and—implicitly, what the geographical distribution of health and education spending is, since total enrollment and health user information can be derived. Comparing this to the distribution of the poor and to the population gives information on:

- the geographical distribution of government spending and necessary reform;
- the potential focus, and locus, of interventions.

If different DHS surveys exist over time, a potentially rich analysis can compare developments in time ion the various indicators. Such comparisons could include how outcome indicators such as education and health evolved, whether backward regions caught up in basic service access and health and education spending; whether the distribution of poverty changed significantly; and whether incidence of education or health spending improved.

For further reference, see http://www.worldbank.org/poverty/health/data/index.htm, Filmer and Pritchett (1999), Gwatkin et al. (2000).

Guide to Web Resources

Web Site	Description
http://www.census.gov/main/www/stat_int.html	United States Census Bureau: List and links to statistical agencies worldwide. These provide information on the latest census, household surveys, and specialized data sets.
http://afr.worldbank.org/aft2/cwiq/overvw.htm	Core Welfare Indicator Questionnaire —joint initiative by World Bank, UNDP, and UNICEF to monitor social indicators in Africa
http://www.worldbank.org/research/journals/wbromast.htm	World Bank *Research Observer*, master list of articles
http://www.worldbank.org/poverty/inequal/methods/index.htm	World Bank Web site on Inequality: Measurement and Decomposition
http://www.worlbank.org/lsms	World Bank Web site on Living Standards Measurement Study—a household survey in measuring and understanding poverty
http://www.macroint.com/dhs	Demographic and health surveys—complete list of surveys available and description of data. Statistics on population, health, and nutrition in developing countries.
http://www.sws.org.ph/swr.htm	Philippine social weather surveys—data on Philippine economic and social conditions
http://www.cgiar.org/ifpri/index.htm.	Web site of International Food Policy Research Institute (IFPRI), a member of the Consultative Group on International Agricultural Research (CGIAR)
http://www.worldbank.org/poverty/inequal/povmap/index.htm	World Bank Web site on Geographic Aspects of Inequality and Poverty
http://www.grida.no/prog/global/poverty/index.htm	Web site of United Nations Environment Program's Global Resource Information Database, use of GIS for agricultural research and poverty mapping