Background on Poverty Mapping

This paper focuses on experience with maps depicting indicators of poverty at a subnational scale, such as by district or community within a given country. Figure 1 provides an example of poverty maps for Ecuador at different levels of resolution, showing the percentage of individuals classified as poor (the so-called head count index).

Figure 1. Poverty maps for Ecuador at various levels of spatial resolution

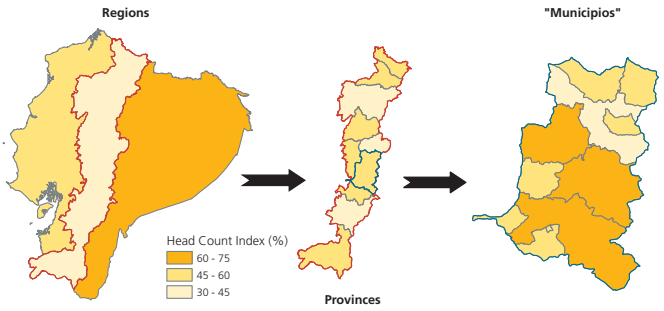

Sources: Boundary files from Centro Internacional de Agricultura (CIAT). Poverty estimates from Hentschel et al. 2000.

High-resolution maps can help uncover poor areas that might otherwise go undetected. Shown here are poverty maps for Ecuador at increasing levels of resolution, from national to regional, provincial, and *municipios* (districts). Higher-resolution maps reveal that the low-poverty region (the Andean region, shown in pale yellow, with a poverty rate of 30-45%) contains several provinces with a considerably higher incidence of poverty (45-60%, shown in dark yellow). One such province, shown at highest resolution, encompasses several *municipios* (districts) with extremely high poverty levels (60-75%, shown in orange) alongside areas of moderate and low poverty.

Higher-resolution maps are useful to decision-makers and researchers in part because they powerfully illustrate the spatial heterogeneity of poverty within a country. They are of special interest to environmental scientists and other researchers working with spatial information on land cover change, ecosystem goods and services, infrastructure development, and market integration, and similar topics with locational aspects.

There is as yet no standard methodology for producing high-resolution poverty maps. Various methods have been used and refinements of technique continue to be developed. (See Box 1 for a description of eight generic steps for producing poverty maps.) Each method has its own particular strengths and weaknesses. Data needs differ depending on the analytical methods chosen, and various methods have different implications for the timeframe and costs involved in conducting the analysis. Moreover, some methods require a higher level of statistical and econometric expertise than do others.

Some of the most commonly used methods of poverty mapping are outlined below (and described in more detail in Box 2). The choice of methods and data sources for poverty mapping should be determined according to the purpose for which the resulting map will be used, which often dictates the appropriate level of precision and resolution. In developing countries, it is also important to take into account the prevailing level of technical and human capacity development.

Preparation of a poverty map may be driven by demand (e.g., need for information and analysis for program design and/or implementation) or by supply (e.g., researcher interest in testing or refining a methodology). Ideally, a poverty mapping exercise will emerge from and be shaped by the process of policy dialogue between map producers and users. Through policy dialogue, map producers and users can work together to explore the specific purposes of a proposed poverty mapping effort. Technical experts can help increase decision-makers' awareness of the potential uses of poverty mapping as well as the inherent limitations of these techniques. Such discussions can help illuminate important issues, not only with respect to choice of method and data source(s), but also spark ideas concerning collaboration between various researchers and institutions, capacity development, dissemination of resulting data products, and long-term sustainability of the mapping effort.

Small area estimation. Poverty maps based on the small area-estimation method rely on sophisticated econometric techniques and a set of identical variables (e.g., household characteristics and educational background) in both a census and a surveyed representative sample of the overall population. By combining census and household survey data, researchers benefit from the strengths of each instrument: a census' complete coverage of a country and a survey's more detailed information. The survey provides the specific poverty indicator and the parameters, based on regression models, to predict the poverty measure for the census.

Box 1 Poverty mapping – generic steps

These eight generic steps involved in a poverty mapping effort highlight key decision points faced by researchers and map producers. Not every poverty-mapping exercise will include all eight steps or follow them sequentially.

1. Define purpose and expected use of mapping

In an ideal world, all poverty mapping would start here. Maps may be needed to show that certain regions are disadvantaged, to rapidly assess options for food emergency interventions, to target public investment to areas of greatest need, or to investigate specific causes of poverty. The purpose and intended use of poverty maps determine the scope and the required precision of the mapping exercise and should shape methodological choices described below.

2. Select measure(s) of poverty and human wellbeing

Choosing an indicator or indicators of poverty is a pivotal step in map production. Poverty is a multidimensional phenomenon, including economic, social, and other aspects of human wellbeing. The selected indicator may be a monetary or non-monetary variable—for example, the proportion of households below a certain income level or the proportion of households without access to sanitation. Researchers sometimes distinguish between status and outcome variables—e.g., access to safe drinking water (status) versus incidence of waterborne diseases (outcome)— but because indicators of poverty are interdependent, the distinction between status and outcome measures is not always clear. A poverty indicator may measure a single important dimension of human wellbeing, such as household expenditure compared to a minimum necessary level or poverty line. Alternatively, the indicator may be multidimensional, for instance, a composite index that depicts deficits in basic human needs, such as education, health care, and sanitation. Each type of poverty indicator has its own strengths and weaknesses, and the choice of indicator will certainly influence who is classified as "poor."

3. Select input data

Data used to construct a poverty map typically are drawn from population or agricultural censuses, household surveys, or spatial (GIS) databases in which values are fixed to specific locations on a grid. Increasingly, poverty mapping relies on data from many sources. Data used in poverty mapping may vary in coverage, collection method, and level of resolution, all of which may have methodological implications. Data coverage may be comprehensive—such as a national census or a detailed map covering the entire geographic area under consideration—or it may be partial, for example, a survey of household expenditures covering a representative sample of the population. Researchers may face choices with respect to data collection methods, including qualitative versus quantitative approaches, or top-down versus participatory methods. The level of resolution of input data used in poverty mapping may be high (e.g., household level) or relatively coarse (e.g., averages for census tracts or administrative units).

4. Select method of estimating or calculating poverty indicator

Researchers may choose to estimate a single variable, such as per capita household expenditures compared to a specific standard of living (i.e., poverty line). Alternatively, they could use a composite index, which may be calculated by simple aggregation (i.e., equal weighting) of a few variables or by multivariate analysis, such as principal components or factor analysis.

Typically, the poverty indicator is an expenditure-based indicator of welfare, such as the proportion of households that falls below a certain expenditure level (i.e., poverty line). In recent years, researchers have relied on two principal methods for their small-area-based poverty maps. The first requires access to detailed household-unit-level data from a census. If such household-unit data are unavailable, unreliable, or incomplete—as is frequently the case in many developing countries—researchers have applied average values for a given indicator at the community level (see Box 2 for more detail).

Small area estimation-based household-level survey data generally are more accurate and reliable than those based on community-based averages. Indeed, the small-area estimation technique using household-unit data is the only poverty mapping method that generates an estimate of statistical error. However, the technical and data requirements of this technique are relatively rigorous, and the approach works best in countries with regular and comprehensive national censuses and household surveys. Community-level averages are more readily available, but using the small-area estimation technique with such data generates an uncertain error, and the datasets used may not provide a good proxy for the poverty indicator that the researcher seeks to measure.

Other poverty mapping methods. Although the "newest" type of poverty maps are based on small-area estimation techniques, other methods have a longer history of application and important lessons have been learned in the course of their use. Many such methods feature the use of composite indexes, including the Human Development Index (HDI) originated by the United Nations Development Programme (UNDP), as well as various basic needs measures. The latter, sometimes

Box 1 continued

5. Select a method to calculate, estimate, or display poverty indicator for geographic area

Depending on the chosen poverty indicator, input data, and method of estimation/calculation, researchers will have different options for calculating or estimating the poverty indicator across a geographic area. For instance, if map producers are using census-level data made available at the household level, then simple aggregation of the data for the selected geographic unit may suffice. However, researchers often need techniques that are more sophisticated. Poverty maps often combine census data (featuring complete country coverage) with household survey data (encompassing a representative sample of the selected population). This is accomplished by means of advanced statistical methods based on econometric techniques, sometimes referred to as *small area estimation*. Combining data from these two sources enables a poverty mapping study to benefit from both the complete spatial coverage of the census and from a relevant poverty indicator in the household survey. Such statistical techniques help overcome the survey's insufficient sample size, which could not be aggregated to small administrative units, and the census' lack of an appropriate poverty measure.

6. Decide on number of units for final map (resolution) to present poverty data

For many poverty-mapping methods, this step is often combined with the previous one. In the case of small area estimation relying on household-unit data, researchers cannot map an individual household; they must aggregate household-level data to larger units to reduce the statistical error in their prediction model. Sensitivity tests conducted by researchers suggest that a minimum of 5,000 households is needed to reduce statistical error to an acceptable level (Elbers et al. 2002). The number of households required may be significantly higher in other cases, especially if the statistical model is not as strong in its predictive power.

7. Produce and distribute maps

Mapping software is used to produce a spatial representation of the geographic distribution of calculated/estimated poverty indicators. Maps and supporting analyses are distributed to the targeted decision-makers. Increasingly, map producers are supplementing hardcopy maps with other products, such as interactive decision-support tools and/or datasets on compact discs (CDs), aimed at various audiences (technical, general, or mixed).

8. Monitor usage and feedback

Poverty maps are used for various purposes, ranging from identifying and understanding the causes of poverty, to assisting in program development and policy formulation, to guiding allocation of anti-poverty investments and expenditures. Map producers should monitor and evaluate the various ways in which their maps are being used by decision-makers and/or researchers, and users should provide feedback on the impact and limitations of poverty maps to map developers.

Sources: Adapted from Henninger (1998), Deichmann (1999), and Davis (2002)

referred to as "unsatisfied basic needs" indexes, have been used primarily in Latin America.

One advantage of composite indexes is that they are intuitive and easy for a general audience to understand. Moreover, this approach requires less advanced statistical expertise than small area estimation. Composite indicators are stronger on the social dimensions of poverty and, on first impression, they appear to better capture the multidimensional nature of human wellbeing.

The most serious criticism of composite indexes is that their weighting of variables can be arbitrary and theoretically unsound. Even a small change in the weighting scheme could easily lead to a change in the proportion of households classified as poor and overturn the ranking of geographic areas identified as poor.

Caveats. Although poverty mapping can be a powerful tool for analyzing poverty and communicating the results to technical and non-technical audiences, experts hasten to point out the limitations of these techniques. Poverty maps are not a panacea for understanding or solving poverty problems; they are only one tool among many for investigating the complex phenomenon of poverty. They should be used in conjunction with other information and analysis that provide context and groundtruthing within communities.

Poverty maps can be used to explore the spatial aspects of various components of human poverty. However, indirect estimation of poverty, as opposed to direct observation in the field, introduces some degree of uncertainty. Careful additional analyses are needed before conclusions are drawn on any meaningful correlation, much less causal relationships, between these variables.

Box 2 Methods of poverty mapping

Expenditure-based small area estimation using household-unit data

This approach was initiated by researchers at the World Bank in 1996 (Hentschel and Lanjouw 1996). The techniques have been further refined, mostly under the leadership of individuals at the World Bank, universities, and in-country partner institutions (e.g., Hentschel et al. 1998, Hentschel et al. 2000, Statistics SA 2000, Alderman et al. 2001, and Elbers et al. 2002). A group of these researchers is currently developing a handbook that describes this approach step by step (Lanjouw 2002).

Typically, this approach begins with a nationally representative household survey, such as the Living Standard Measurement Survey, to acquire a reliable estimate of household expenditure (y) and calculate more specific poverty measures linked to a poverty line. A common set of explanatory variables x (e.g., educational background, household characteristics, and quality of housing) at the household-unit level in both the survey and the census is then used to estimate the statistical relationship between y and x in the survey. Once a robust model has been identified for the survey, researchers apply the final model to the census data at household-unit level to predict per capita household expenditures (including an error estimate). These household-unit data can then be aggregated to small statistical areas, such as districts, to obtain more robust estimates of the percentage of households living below the poverty line. Finally, these poverty rates by administrative area are linked to a mapping program to produce a poverty map showing the spatial distribution of poverty. In most cases, the spatial resolution of this map, i.e., the number of administrative units, is significantly higher than would be possible using the household survey alone.

This small area approach using household-unit data has been applied in various countries, including Ecuador, Guatemala, Nicaragua, Panama, and South Africa. It is currently being considered for map development in a number of other countries, including China, Indonesia, Kazakhstan, Kenya, Kyrgyzstan, Madagascar, Malawi, Mexico, Mozambique, Pakistan, Thailand, and Uganda.

In all of the examples cited above, the outcome variable has been an estimate of household expenditure linked to a poverty line. In principle, this small-area estimation technique could be applied using a different outcome variable, for example, a non-monetary indicator; however, no example of such an approach has yet been published. As this report was going to press, a small area estimation of Demographic and Health Survey (DHS) indicators was initiated for Cambodia (Montana 2002). Similarly, researchers at the International Food Policy Research Institute (IFPRI) and Cornell University started on a small-area estimation that intends to use nutrition indicators as outcome variables for Tanzania and Ecuador, respectively (Minot 2002).

Davis (2002) lists the following strengths of this approach: It is relatively easy for national analysts familiar with econometric modeling to check the reliability of their estimates, because the data processing program provided by the World Bank is equipped with an error estimation module. It is the only method "where statistical properties have been—and continue to be—thoroughly investigated." In addition, the method has institutional support from the World Bank and a team of researchers is available to further refine the method and provide technical assistance.

One limitation of this approach may be that it is less feasible for individual researchers who would like to work independently. Without institutional support from the World Bank and a collaborative research agreement, an independent researcher may not be able to obtain access to household-unit data. A second important limitation is that census data in many developing countries may not provide a sufficient number of explanatory variables to build a robust statistical model. It is no surprise that most of the countries where this small area estimation has been used have a regular and comprehensive effort of national censuses and household surveys, with relatively strong statistical departments and relevant statistical expertise. Other obstacles—such as enormous input data files, literally millions of records to be processed, and the not-trivial econometric expertise required—are becoming less of an issue because of the growing processing power of personal computers as well as the technical support provided by the World Bank.

In addition, it is important that poverty mapping is always seen in the overall context of a country's decision-making processes. Technical tools like poverty maps run the risk of being abandoned once initial donor support has waned. To ensure a path of sustained use and support for poverty maps, fundamental questions need to be addressed, such as how to retain skilled analysts in the public sector, overcome limited or lacking demand and funding from policymakers, and convince decision-makers that continued investment in poverty maps is worthwhile in an environment that does not follow a purely technical approach to decision-making.

Box 2 *continued*

Expenditure-based small area estimation using community-level data

This approach has been pioneered by researchers at the World Bank and centers within the CGIAR system (Minot 1998, Bigman et al. 2000, Bigman and Fofack 2000, and Minot 2000). While different researchers may have varied their specific approaches—for example, by working at a regional versus a village level—all methods have a common element. The input data for the statistical model relies on average values—for example, for communities or regions—and not on household-unit data. For example, Bigman et al. (2000) combined such data from household surveys and a census, and also used variables generated by a GIS (i.e., distance to schools, water points, etc.) to overcome constraints on the quality and availability of household-level census data in Burkina Faso.

Besides this main difference, the small area estimation follows an approach similar to the one outlined above. First, researchers determine the statistical relationship between household expenditure and a set of explanatory variables within the survey data alone. Once they have determined a strong relationship between the two sets of variables, they can apply this relationship to the same variables in the census, but this time for a community or an administrative region, not an individual household.

This approach, using community-level averages, has been applied in Burkina Faso and Vietnam. Other examples, not examined here, have been reported for Kenya (Bigman and Loevinsohn 1999) and India (Bigman and Srinivasan 2001).

Davis (2002) mentions that more readily available data is a plus for this approach. Because of legitimate concerns about data confidentiality, government agencies are more willing to provide researchers with census data on community averages than on households. He also points out two limitations: In some cases, the averages calculated for the community or the region may not be a good proxy for the distribution of poverty. Second and most importantly, the error associated with such an estimation approach has not been thoroughly investigated yet. It is not clear how much statistical reliability is sacrificed for data access, and what the most appropriate use for this approach is. Differences between the community-level and the household-unit approach need to be systematically investigated. A first paper examining how the levels of precision differ between the two approaches was presented in 2002 (Minot and Baulch 2002a).

Other methods not based on small area estimation

Some poverty mapping techniques use composite indexes as the poverty measure and rely on the direct aggregation of census data to display the poverty indicator for the chosen geographic area. Please refer to Davis (2002) or Henninger (1998) for detailed descriptions and examples of other poverty mapping approaches that do not rely on small area estimation.

Composite indexes used for poverty mapping studies include UNDP's well known Human Development Index and various basic needs measures, sometimes also referred to as "unsatisfied basic needs" index. Basic needs indexes have been applied primarily in Latin America. The Human Development Index is based on three variables: life expectancy, education (literacy), and income. All components are weighted equally. Basic needs indexes typically have included more than three variables—for example, literacy, access to water, access to sanitation, access to health services, and quality of housing. Many of the existing basic needs indexes have equal weighting schemes similar to the HDI. Others have relied on expert opinion or multivariate statistical techniques to provide weightings for each variable.

Sources: Adapted from Henninger (1998), Deichmann (1999), and Davis (2002).