



United Nations
University

WIDER

World Institute for Development Economics Research

Discussion Paper No. 2002/39
**Producing an Improved Geographic
Profile of Poverty**

Methodology and Evidence from Three
Developing Countries

Gabriel Demombynes,¹ Chris Elbers,²
Jenny Lanjouw,³ Peter Lanjouw,⁴
Johan Mistiaen⁵ and Berk Özler⁶

March 2002

Abstract

This paper implements a methodology for estimating poverty in Ecuador, Madagascar and South Africa, at levels of disaggregation that to date have not generally been available. The methodology is based on a statistical procedure to combine household survey data with population census data, imputing into the latter a measure of per capita consumption from the former. The countries are very unlike each other—with different geographies, stages of development, quality and types of data, and so on. Yet the paper demonstrates that in all three countries the poverty estimates produced from census data are both *plausible* (in that they match well stratum-level estimates calculated directly from the household surveys) and satisfactorily *precise* (at a level of disaggregation far below that allowed by household surveys). .../...

Keywords: poverty measurement, poverty profiles, spatial distribution, forecasting models, statistical inference

JEL classification: I32, O18, C53

Copyright © author(s) 2002

¹University of California, Berkeley and World Bank; ²Vrije (Free) University, Amsterdam; ³Yale University and Brookings Institution; ⁴World Bank; ⁵University of Maryland and World Bank; ⁶World Bank

This study is linked to the theme of Growth and Poverty discussed at the WIDER Development Conference, 25-26 May 2001, Helsinki. UNU/WIDER gratefully acknowledges the financial support from the governments of Denmark, Finland and Norway to the 2000-01 research programmes.

The paper illustrates how the resulting poverty estimates can be represented in maps, thereby conveying much information about the magnitude of poverty across localities, as well as the precision of estimates, in a way which can be readily absorbed by non-technical audiences. The paper finally notes that perceptions as to the importance of geographical dimensions of poverty are themselves a function of the degree of spatial disaggregation of available estimates of poverty. The smaller the localities into which a country can be broken down the more likely one will conclude that geography matters.

Acknowledgements

We are grateful to the Instituto Nacional de Estadística y Censo (INEC), Ecuador, Statistics South Africa, and Institut National de la Statistique (INSTAT), Madagascar, for access to their unit record census data and the Bank Netherlands Partnership Program (BNPP) for financial support. Useful comments were received from Jesko Hentschel, Jose Lopez-Calix, and Martin Ravallion. The views in this paper are our own and should not be taken to reflect those of the World Bank or any of its affiliates. All errors are our own.

UNU World Institute for Development Economics Research (UNU/WIDER) was established by the United Nations University as its first research and training centre and started work in Helsinki, Finland in 1985. The purpose of the Institute is to undertake applied research and policy analysis on structural changes affecting the developing and transitional economies, to provide a forum for the advocacy of policies leading to robust, equitable and environmentally sustainable growth, and to promote capacity strengthening and training in the field of economic and social policy making. Its work is carried out by staff researchers and visiting scholars in Helsinki and through networks of collaborating scholars and institutions around the world.

UNU World Institute for Development Economics Research (UNU/WIDER)
Katajanokanlaituri 6 B, 00160 Helsinki, Finland

Camera-ready typescript prepared by Lorraine Telfer-Taivainen at UNU/WIDER
Printed at UNU/WIDER, Helsinki

The views expressed in this publication are those of the author(s). Publication does not imply endorsement by the Institute or the United Nations University, nor by the programme/project sponsors, of any of the views expressed.

ISSN 1609-5774

ISBN 92-9190-200-4 (printed publication)

ISBN 92-9190-201-2 (internet publication)

1 Introduction

Poverty maps provide a detailed description of the spatial distribution of poverty. Detailed geographic profiles of poverty can be extremely valuable to governments, nongovernmental organizations and multilateral institutions that want to strengthen the impact that their spending has on poverty. For example, many developing countries use poverty maps to guide the division of resources among local agencies or administrations as a first step in reaching the poor.

Poverty maps can also be an important tool for research. Recent theoretical advances have brought income and wealth distribution back into a prominent position in growth and development theories.¹ Distributions of wellbeing are also considered determinants of specific socioeconomic outcomes, such as individual health or levels of violence.²

Construction of detailed geographic poverty profiles and empirical testing of the importance of theoretical relationships, however, has been held back by the poor quality of distributional data. The problem is that the detailed household surveys which include reasonable measures of income or consumption are samples, and are rarely representative or of sufficient size at low levels of disaggregation to yield statistically reliable estimates. In the three developing countries that we examine in this paper the lowest level of disaggregation possible using sample data is to regions, which encompass hundreds of thousands of households. At the same time, census (or other large sample) data, which are of sufficient size to allow disaggregation, either have no information about income or consumption, or measure these variables poorly.³

This paper describes a recently developed statistical procedure to combine household and census data, which takes advantage of the detailed information available in household sample surveys and the comprehensive coverage of a census (Elbers, Lanjouw and Lanjouw, 2001). Using a household survey to impute missing information in the census we estimate (as opposed to directly measure) poverty and inequality at a disaggregated level based on a household per capita measure of expenditure, y_h . The idea is straightforward.

¹ See for example, Murphy, Shleifer and Vishny (1989), Galor and Zeira (1993); Banerjee and Newman (1993); Aghion and Bolton (1997); Alesina and Rodrik (1994); Persson and Tabellini (1994) for early contributions to this rapidly growing literature.

² Deaton (1999) argues that it is most reasonable to search for a relationship between individual health outcomes and local, rather than national, income inequality. Demombynes and Özler (2001) explore the relationship between local inequality and crime in South Africa.

³ For example, a single question regarding individuals' incomes in the 1996 South African census generates an estimate of national income just 83 percent the size of the national *expenditure* estimate derived from a representative household survey, and a per capita poverty rate 25 percent higher, with discrepancies systematically related to characteristics such as household location (Alderman et al., 2000). In Brazil the PNAD household survey, covering a very large sample, is thought to yield an unreliable measure of household income (see Elbers, Lanjouw, Lanjouw and Leite, 2001).

First a model of y_h is estimated using the sample survey data, restricting explanatory variables to those that can be linked to households in both sets of data. Then, letting W represent an indicator of poverty or inequality, we estimate the expected level of W given the census-based observable characteristics of the population of interest using parameter estimates from the ‘first stage’ model of y . The same approach could be used with other household measures of wellbeing, such as per capita expenditure adjusted by equivalences scales, or to estimate inequalities in the distribution of household characteristics other than expenditure, such as assets, income, or employment. A recent study using data from Brazil extends the approach to combine a detailed but small sample survey with a much larger sample survey dataset rather than the full unit record level census (Elbers, Lanjouw, Lanjouw and Leite, 2001).

Drawing on evidence from three different countries—Ecuador, Madagascar and South Africa—we illustrate that the method generates reliable estimates of poverty at a very disaggregated level. Our estimates, for instance, of headcount rates of poverty for ‘counties’ of around 1000-2000 households have 95 percent confidence intervals approximately the same size as those of stratum (region) level estimates in the household surveys. With good welfare estimates for groups the size of towns, villages or even neighbourhoods, policymakers have a valuable tool for targeting purposes, and researchers are able to test a variety of hypotheses at levels of disaggregation where assumptions about stable underlying structures are more tenable than at a cross-country level. That the method performs satisfactorily in three settings as dissimilar as the countries considered in this paper lends support to the notion that the approach will be useful in many contexts. However, it is important to emphasize that data requirements are non-negligible and unlikely to be satisfied everywhere at once.

2 An overview of the methodology

The basic methodology is broadly straightforward. First, the survey data are used to estimate a prediction model for either consumption or incomes. The selection of exogenous variables is restricted to those variables that can also be found in the census (or some other large dataset) or in a tertiary dataset that can be linked to both the census and survey. The parameter estimates are then applied to the census data and poverty statistics derived. The key assumption is that the models estimated from the survey data apply to census observations. This is most reasonable if the survey and census years coincide. If different years are used but the assumption is considered reasonable, then the welfare estimates obtained refer to the census year, whose explanatory variables form the basis of the predicted expenditure distribution.

Simple checks can be carried out to compare basic poverty or inequality statistics across the two datasets. An important feature of the approach applied here involves the explicit recognition that the poverty or inequality statistics estimated using a model of income or

consumption are statistically imprecise. Standard errors must be calculated. The following subsections briefly summarize the discussion in Elbers, Lanjouw and Lanjouw (2001).

2.1 Definitions

Per capita household expenditure, y_h , is related to a set of observable characteristics, \mathbf{x}_h , that can be linked to households in both the household survey and the census:⁴

$$\ln y_h = E[\ln y_h | \mathbf{x}_h] + u_h \quad (1)$$

Using a linear approximation to the conditional expectation, we model the observed log per capita expenditure for household h as:

$$\ln y_h = \mathbf{x}_h \boldsymbol{\beta} + u_h \quad (2)$$

where $\boldsymbol{\beta}$ is a vector of k parameters and u_h is a disturbance term satisfying $E[u_h | \mathbf{x}_h] = 0$. In applications we allow for location effects and heteroskedasticity in the distribution of the disturbances.

The model in (2) is estimated using the household survey data. We are interested in using these estimates to calculate the welfare of an area or group for which we do not have any, or insufficient, expenditure information. Although the disaggregation may be along any dimension—not necessarily geographic—for convenience we will refer to our target population as a ‘county’. Household h has m_h family members. While the unit of observation for expenditure in these data is the household, we are more often interested in poverty and inequality measures based on individuals. Thus we write $W(\mathbf{m}_v, \mathbf{X}_v, \boldsymbol{\beta}, \mathbf{u}_v)$, where \mathbf{m} is a vector of household sizes, \mathbf{X} is a matrix of observable characteristics and \mathbf{u} is a vector of disturbances. Because the disturbances for households in the target population are always unknown, we consider estimating the expected value of the indicator given the census households’ observable characteristics and the model of expenditure in (2).⁵ We denote this expectation as:

$$\mu_v = E[W | \mathbf{m}_v, \mathbf{X}_v, \boldsymbol{\xi}] \quad (3)$$

where $\boldsymbol{\xi}$ is the vector of model parameters, including those which describe the distribution of the disturbances. In constructing an estimator of μ_v we replace the unknown vector $\boldsymbol{\xi}$ with consistent estimators, $\hat{\boldsymbol{\xi}}$, from the first stage expenditure regression. This yields

⁴ The explanatory variables are observed values and need to have the same degree of accuracy in addition to the same definitions across data sources. From the point of view of our methodology it does not matter whether these variables are exogeneous.

⁵ If the target population includes sample survey households then some disturbances are known. As a practical matter we do not use these few pieces of direct information on y .

$\hat{\mu}_v = E[W | \mathbf{m}_v, \mathbf{X}_v, \hat{\xi}]$. This expectation is generally analytically intractable so we use Monte Carlo simulation to obtain our estimator, $\tilde{\mu}_v$.

2.2 Properties

The difference between $\tilde{\mu}_v$, our estimator of the expected value of W for the county, and the *actual* level of welfare for the county may be written (suppressing the index v):

$$W - \tilde{\mu} = (W - \mu) + (\mu - \hat{\mu}) + (\hat{\mu} - \tilde{\mu}) \quad (4)$$

Thus the prediction error has three components: the first due to the presence of a disturbance term in the first stage model which implies that households' actual expenditures deviate from their expected values (idiosyncratic error); the second due to variance in the first stage estimates of the parameters of the expenditure model (model error); and the third due to using an inexact method to compute $\hat{\mu}$ (computation error).⁶

Idiosyncratic error

The variance in our estimator due to idiosyncratic error V_I falls approximately proportionately in the size of the population of households in the county. In other words, the smaller the target population, the greater is this component of the prediction error, and there is thus a practical limit to the degree of disaggregation possible. At what population size this error becomes unacceptably large depends on the explanatory power of the \mathbf{x} variables in the expenditure model and, correspondingly, the importance of the remaining idiosyncratic component of the expenditure.

Model error

To assess the variance due to model error V_M we can employ the delta method: $V_M \approx \nabla^T V(\hat{\xi}) \nabla$, where $\nabla = [\partial \tilde{\mu} / \partial \xi]_{\hat{\xi}}$ and $V(\hat{\xi})$ is the asymptotic variance covariance matrix of the first stage parameter estimators. Because this component of the prediction error is determined by the properties of the first stage estimators, it does not increase or fall systematically as the size of the target population changes. Its magnitude depends, in general, only on the precision of the first stage coefficients and the sensitivity of the indicator to deviations in household expenditure. For a given county its magnitude will also depend on the distance of the explanatory variables for households in that county from the levels of those variables in the sample data.

Computation error

The variance in our estimator due to computation error depends on the method of computation used. As our calculations of the idiosyncratic and models errors are based on

⁶ Elbers, Lanjouw, Lanjouw, and Leite (2001) use two surveys, rather than a survey and census, which then also introduces sampling error.

simulations, we can make the computation error become as small as desired by choosing a large enough number of simulation draws (at the cost of computational resources and time).

We use Monte Carlo simulations to calculate: $\hat{\mu}$, the expected value of the poverty r inequality measure conditional on the first stage model of expenditure; V_r the variance in $\hat{\mu}$ due to the idiosyncratic component of household expenditures; and, for use in determining the model variance, the gradient vector $\nabla = [\partial \tilde{\mu} / \partial \xi]_{|\xi}$. Let the vector $\hat{\mathbf{u}}^r$ be the r th draw from our estimated disturbance distribution. With each vector of simulated disturbances we construct a value for the indicator, $\hat{W}_r = W(\mathbf{m}, \mathbf{X}, \hat{\xi}, \hat{\mathbf{u}}^r)$, where \mathbf{m} and \mathbf{X} represent numbers of people and observable characteristics of census households, respectively. The simulated expected value for the indicator is the mean over R replications:

$$\tilde{\mu} = \frac{1}{R} \sum_{r=1}^R \hat{W}_r \quad (5)$$

The variance of W around its expected value due to the idiosyncratic component of expenditures can be estimated in a straightforward manner using the same simulated values:

$$\hat{V}_r = \frac{1}{R} \sum_{r=1}^R (\hat{W}_r - \tilde{\mu})^2 \quad (6)$$

Simulated numerical gradient estimators are constructed as follows: We make a positive perturbation to a parameter estimate, say $\hat{\beta}_k$, by adding $\delta |\hat{\beta}_k|$, and then calculate $\tilde{\mu}^+$. A negative perturbation of the same size is used to obtain $\tilde{\mu}^-$. The simulated central distance estimator of the derivative $\partial \tilde{\mu} / \partial \beta_k |_{\xi}$ is $(\tilde{\mu}^+ - \tilde{\mu}^-) / (2\delta |\hat{\beta}_k|)$. Having thus derived an estimate of the gradient vector, we can calculate $\hat{V}_M = \nabla^T V(\hat{\xi}) \nabla$.

3 Data

In all three of the countries examined here, household survey data were combined with unit record census data. In Ecuador the poverty map is based on census data from 1990, collected by the National Statistical Institute of Ecuador (Instituto Nacional de Estadística y Censos, INEC) combined with household survey data from 1994. The census covered roughly 2 million households. The sample survey (Encuesta de Condiciones de Vida, ECV) is based on the Living Standards Measurement Surveys approach developed by the World Bank, and covers just under 4,500 households. The survey provides detailed information on a wide range of topics; including food consumption, nonfood consumption, labor activities, agricultural practices, entrepreneurial activities, and access to services such as education and health. The survey design incorporates both clustering and stratification

on the basis of the country's three main agroclimatic zones and a rural-urban breakdown. It also oversamples Ecuador's two main cities, Quito and Guayaquil. Hentschel and Lanjouw (1996) develop a consumption aggregate for each household, and also adjust these for spatial price variation based on a Laspeyres food price index reflecting the consumption patterns of the poor. World Bank (1996) developed a consumption poverty line of 45,476 sucres per person per fortnight (approximately \$1.50 per person per day) which underpins the poverty numbers reported here. It is important to recognize that the 1994 ECV data were collected four years after the census, but that the methodology described above is predicated on the model of consumption in 1994 being appropriate for 1990. Because the period 1990-4 was one of relative stability in Ecuador, it is not unreasonable to assume relatively little change in the underlying model over this time period. Comparative summary statistics on a selection of common variables from the two data sources support the presumption of little change over the period. Details on these data and application of the poverty mapping methodology in Ecuador can be found in Hentschel, Lanjouw, Lanjouw, and Poggi (2000) and Elbers, Lanjouw, Lanjouw and Leite (2001).

Three data sources were used to produce local level poverty estimates for Madagascar. First, the 1993 unit record population census data were collected by the Direction de la Démographie et Statistique Social (DDSS) of the Institut National de la Statistique (INSTAT). Second, a household survey, the Enquête Permanente Auprès des Ménages (EPM) was fielded to over 4,508 households between May 1993 and April 1994, by the Direction des Statistique des Ménages (DSM) of INSTAT. Third, we made use of a variety of spatial and environmental outcomes at the Fivondrona level (e.g. representing a collection of Firaisanas or 'communes'). These data were specifically provided to this project by CARE. The household level welfare indicator underpinning the Madagascar poverty map includes components such as an imputed stream of consumption from the ownership of consumer durables, so as to be as comprehensive as possible. Further details about the analysis in Madagascar are provided in Mistiaen et al. (2001).

Three datasets also underpin the South African poverty map. The first is the OHS (October Household Survey), an annual survey which focuses on some key indicators of living-patterns in South Africa. In particular it focuses on employment, internal migration, housing, access to services, individual education, and vital statistics. 29,700 households were interviewed in the 1995 round of the survey. The IES (Income and Expenditure Survey) is the second source of data, providing information on the income and expenditure of households for the 12-month period prior to the interview. The IES was designed for use with the OHS. While the interviews for the IES were conducted at a slightly later date than the OHS, the same households were visited. In all, 28,710 households remained in the dataset after the two surveys were merged. The third source of data, the population census of 1996, covers over 8.3 million households. In addition to information on household composition it collected some details on housing and services in a manner that paralleled the OHS. Further details on the South African data and analyses can be found in Alderman et al. (2001).

4 Implementation

In all three countries implementation follows a broadly similar procedure. The first stage estimation is carried out using the household sample survey. For each of the three countries considered in this paper, the respective household survey is stratified into a number of regions and is representative at the level of those regions. Within each region there are one or more levels of clustering. At the final level, households are randomly selected from a census enumeration area. Such groups we refer to as ‘cluster’ and denote by a subscript c . Expansion factors, l_{ch} , allow the calculation of regional totals.

Our first concern is to develop an accurate empirical model of household consumption. Consider the following model:

$$\ln y_{ch} = E[\ln y_{ch} | x_{ch}^T] + u_{ch} = x_{ch}^T \beta + \eta_c + \varepsilon_{ch} \quad (9)$$

where η and ε are independent of each other and uncorrelated with observables, x_{ch} . This specification allows for an intracluster correlation in the disturbances. One expects location to be related to household income and consumption, and it is certainly plausible that some of the effect of location might remain unexplained even with a rich set of regressors. For any given disturbance variance, σ_{ch}^2 , the greater the fraction due to the common component η_c , the less one enjoys the benefits of aggregating over more households within a county. Welfare estimates become less precise. Further, the greater the part of the disturbance which is common, the lower will be inequality. Thus, failing to take account of spatial correlation in the disturbances could result in underestimated standard errors on all welfare estimates, and upward biased estimates of inequality.

Since unexplained location effects reduce the precision of poverty estimates, the first goal is to explain the variation in consumption due to location as far as possible with the choice and construction of x_{ch} variables. To varying degrees in turn for Ecuador, Madagascar, and South Africa, we try to tackle this in four ways.

1. We estimate different models for each stratum in the country’s respective survey.
2. We include in our specification household level indicators of connection to various networked infrastructure services, such as electricity, piped water, networked waste disposal, telephone etc. To the extent that all or most households within a given neighborhood or community are likely to enjoy similar levels of access to such networked infrastructure, these variables might capture unobserved location effects.
3. Third, we calculate means at the enumeration area (EA) level in the census (generally corresponding to the ‘cluster’ in the household survey) of household level variables, such as the average level of education of household heads per cluster. We then insert these cluster means into the household survey and consider them for inclusion in the

first stage regression specification.⁷ These cluster level variables also serve to proxy location-specific correlates of expenditure.

4. Finally, in the case of Madagascar we have merged the Firaisana level dataset provided by CARE and also consider the spatially referenced environmental variables contained in that dataset for inclusion in our household expenditure models.

We apply a selection criterion when deciding on our final specification, requiring a significance level of five percent of all household level regressors. To select location variables (EA means and for Madagascar, the CARE variables), we estimate a regression of the total residuals, \hat{u} , on cluster fixed effects. We then regress the cluster fixed-effect parameter estimates on our location variables and select those that best explain the variation in the cluster fixed-effects estimates.⁸ These location variables are then added to our household level variables in the first stage regression model.

We apply a Hausman test described in Deaton (1997) to determine whether each regression should be estimated with household weights. \bar{R}^2 's in our models are generally high, ranging between 0.45 and 0.77 in Ecuador, 0.29 to 0.63 in Madagascar, and 0.47 to 0.72 in South Africa.⁹ We next model the variance of the idiosyncratic part of the disturbance, $\sigma_{\varepsilon, ch}^2$. Note that the total first stage residual can be decomposed into uncorrelated components as follows:

$$\hat{u}_{ch} = \hat{u}_c + (\hat{u}_{ch} - \hat{u}_c) = \hat{\eta}_c + e_{ch} \quad (10)$$

where a subscript '.' indicates an average over that index. Thus the mean of the total residuals within a cluster serves as an estimate of that cluster's location effect. To model heteroskedasticity in the household-specific part of the residual, we choose between 10 and 20 variables, \mathbf{z}_{ch} , that best explain variation in e_{ch}^2 out of all potential explanatory variables, their squares, and interactions.¹⁰ We estimate a logistic model of the variance of ε_{ch} conditional on \mathbf{z}_{ch} , bounding the prediction between zero and a maximum A set equal to $(1.05) * \max\{e_{ch}^2\}$:

⁷ In Madagascar the EA in the household survey is not the same as that in the census. The most detailed spatial level at which we can link the two datasets is the Firaisana ('commune'). Thus, only Firaisana-level means of various variables from the census data were merged into the household survey. Also in South Africa, the means are calculated at the magisterial district level rather than cluster.

⁸ To avoid overfitting the stratum level regressions (depending on country, these can include as few as 250 households) a maximum of between 5 to 10 EA mean variables were accommodated in the first stage regressions.

⁹ For reasons of space we do not reproduce here the parameter estimates and full set of diagnostics for all 29 regression models. See Elbers et al. (2001), Mistiaen et al. (2001) and Alderman et al. (2001) for further details.

¹⁰ Once again, we limit the number of explanatory variables to be cautious about overfitting.

$$\ln\left[\frac{e_{ch}^2}{A - e_{ch}^2}\right] = z_{ch}^T \hat{\alpha} + r_{ch} \quad (11)$$

Letting $\exp\{z_{ch}^T \hat{\alpha}\} = B$ and using the delta method, the model implies a household specific variance estimator for ε_{ch} of

$$\hat{\sigma}_{\varepsilon, ch}^2 = \left[\frac{AB}{1+B}\right] + \frac{1}{2} \text{Var}(r) \left[\frac{AB(1-B)}{(1+B)^3}\right] \quad (12)$$

Finally, we check whether η and ε are distributed normally, based on the cluster residuals $\hat{\eta}_c$ and standardized household residuals $e_{ch}^* = \frac{e_{ch}}{\hat{\sigma}_{\varepsilon, ch}} - \left[\frac{1}{H} \sum_{ch} \frac{e_{ch}}{\hat{\sigma}_{\varepsilon, ch}}\right]$, respectively where H is the number of households in the survey. The second term in e_{ch}^* is not needed when first stage regressions are not weighted. In many cases normality is rejected, although the standard normal does occasionally appear to be the better approximation even if formally rejected. Elsewhere we use t distributions with varying degrees of freedom (usually 5), as the better approximation. Before proceeding to simulation, the estimated variance-covariance matrix, $\hat{\Sigma}$, is used to obtain GLS estimates of the first stage parameters, $\hat{\beta}_{GLS}$, and their variance, $\text{Var}(\hat{\beta}_{GLS})$.

In Section 2 we outlined a procedure for calculating standard errors around our estimated poverty rates. That procedure has the attraction of allowing the analyst to not only obtain a measure of the overall variance around a particular point estimate, but also to break down that variance into its idiosyncratic and model subcomponents. For the first stage model specification phase it is useful to be able to scrutinize the error components around the final poverty estimates; in practice it is often necessary to return to the specification phase when overall precision of the point estimate deteriorates as a result of inclusion of a particular regressor in the consumption model or in the heteroskedasticity model.

However, once all regression specification issues have been decided, a more direct approach can be implemented to calculate standard errors on the poverty estimates. This approach calculates just the overall variance around the poverty estimates, but does so much more quickly than the procedure described in Section 2. In this approach, we conduct a series of simulations, and for each simulation we draw a *set* of beta and alpha coefficients, $\tilde{\beta}$ and $\tilde{\alpha}$, from the multivariate normal distributions described by the first stage point estimates and their associated variance-covariance matrices. Additionally, we draw $\tilde{\sigma}_{\eta}^2$, a simulated value of the variance of the location error component.¹¹ Combining the alpha coefficients with census data, for each census household we estimate $\tilde{\sigma}_{\varepsilon, ch}^2$, the household-specific variance of the household error component. Then for each household

¹¹ The $\tilde{\sigma}_{\eta}^2$ value is drawn from a gamma distribution defined so as to have mean $\hat{\sigma}_{\eta}^2$ and variance $\hat{V}(\hat{\sigma}_{\eta}^2)$

we draw simulated disturbance terms, η_c^r and ε_{ch}^r , from their corresponding distributions.¹² We simulate a value of expenditure for each household, \hat{y}_{ch} , based on both predicted log expenditure, $\mathbf{x}'_{ch}\tilde{\boldsymbol{\beta}}$, and the disturbance terms:

$$\hat{y}_{ch}^r = \exp(\mathbf{x}'_{ch}\boldsymbol{\beta}^r + \eta_c^r + \varepsilon_{ch}^r) \quad (13)$$

Finally, the full set of simulated \hat{y}_{ch} values are used to calculate expected values of poverty measures for each ‘county’.¹³ We repeat this procedure 100 times, drawing a new set of coefficients and disturbance term for each simulation. For each county, we take the mean and standard deviation of our poverty measures and average expenditure over all 100 simulations. For any given location, these means constitute our point estimates of the poverty rates and average expenditure, while the standard deviations are the standard errors of these estimates.

5 Results

In this section we examine the success of the approach outlined in previous sections in our three case study settings: Ecuador, Madagascar and South Africa. We begin by examining the degree to which our poverty estimates from the census match sample estimates from the countries’ respective surveys at the level at which that those surveys are representative (usually the stratum). We then ask how far we can disaggregate our census-based poverty estimates, when we take the survey-based sampling errors to indicate acceptable levels of precision. We then turn to the ultimate goal of the analysis, namely, to produce disaggregated spatial profiles of poverty. We illustrate how projecting poverty estimates onto maps produces a quick and appealing way in which to convey a considerable amount of information on the spatial distribution of poverty to users. We also show that conclusions as to the spatial heterogeneity of poverty are a direct function of the degree of disaggregation. This implies that by their very nature, sample surveys are likely to lead analysts to understate the significance of spatial variation in poverty.

5.1 How well do survey and census estimates match?

Tables 1-3 present stratum-level estimates of the poverty headcount in our three countries. Table 1 illustrates that measures of the incidence of poverty in Ecuador at the stratum-level are reasonably close to those from the census. Except for Guayaquil and Rural Sierra, the pairs of poverty estimates are comfortably within each others’ 95 percent confidence

¹² We allow for non-normality in the distribution of both η_c and ε_{ch} . For each distribution, we choose a Student's t-distribution with degrees of freedom such that its kurtosis most closely matches that of our first stage residual components, $\hat{\eta}_c$ or e_{ch} .

¹³ Because we are interested in measures based on *per capita* expenditure, these calculations are performed weighted by household size.

intervals and are close to coinciding in several instances. The differences in estimates for Guayaquil and Rural Sierra can presumably be traced to changes in the exogenous variables underpinning the consumption regressions between the 1990 census and the 1994 household survey.¹⁴

Table 1
Stratum-level poverty rates in Ecuador (headcount)

Stratum	Household Survey (s.e)	Census (s.e)
Rural Costa	0.50 (0.042)	0.501 (0.024)
Urban Costa	0.25 (0.03)	0.258 (0.015)
Guayaquil	0.29 (0.027)	0.380 (0.019)
Rural Sierra	0.43 (0.027)	0.527 (0.019)
Urban Sierra	0.19 (0.026)	0.211 (0.027)
Quito	0.25 (0.033)	0.223 (0.022)
Rural Oriente	0.67 (0.054)	0.590 (0.025)
Urban Oriente	0.20 (0.05)	0.189 (0.021)

Source: ECV (1994); Ecuador Census (1990).

Note: Standard errors on poverty estimates from the household survey reflect a 2-stage sampling design effect. Standard errors on poverty estimates from the census have been calculated according to the procedure outlined in the text.

In Madagascar the data refer to the same period. In this country, the main source of concern is that in one or two of the strata, the explanatory power of the first stage regressions is not particularly high (an adjusted R^2 of 0.292 is the lowest obtained in any of our models and applies to the stratum of rural Antsiranana). The resulting relatively high standard errors on our census level predicted poverty estimates make it difficult to reject that they are the same as the sample estimates. However, for rural Antsiranana the point estimates are close to coinciding. For the other strata, as well, the matching between the census and survey estimates is uniformly close, with in no case point estimates falling outside respective confidence intervals.¹⁵

¹⁴ Other factors may also play a role: changes in the definition of urban/rural, or of metropolitan boundaries; non-sampling errors in data entry and data collection; non-applicability of our maintained assumption that stratum-level regression parameters are applicable for sub-stratum localities; etc.

¹⁵ Mistiaen et al. (2001) document that the figures are similarly close for FGT1 and FGT2 measures of poverty.

South African results are also satisfactory (Table 3). Point estimates across the two data sources match closely at the stratum-level such that we cannot reject equality at a 5 percent significance level. Once again, stratum-level standard errors in the IES survey are not small, despite a sample size which is several times larger than the typical LSMS-style household survey.

Table 2
Stratum-level poverty rates in Madagascar* (headcount)

Stratum	Household Survey (s.e)	Census (s.e)
Antananarivo Urban	.544 (.048)	.456 (.017)
Antananarivo Rural	.767 (.037)	.732 (.015)
Fianarantsoa Urban	.674 (.059)	.695 (.017)
Fianarantsoa Rural	.769 (.049)	.781 (.025)
Taomasina Urban	.599 (.086)	.567 (.024)
Taomasina Rural	.810 (.035)	.774 (.025)
Mahajanga Urban	.329 (.072)	.370 (.036)
Mahajanga Rural	.681 (.065)	.650 (.040)
Toliara Urban	.715 (.086)	.653 (.032)
Toliara Rural	.817 (.042)	.820 (.027)
Antsiranana Urban	.473 (.087)	.345 (.027)
Antsiranana Rural	.613 (.073)	.616 (.045)

Source: EPM (1994); Madagascar Census (1993).

Note: Standard errors on poverty estimates from the household survey reflect a 2-stage sampling design effect. Standard errors on poverty estimates from the census have been calculated according to the procedure outlined in the text. *Madagascar estimates are preliminary and subject to revision, see Mistiaen et al. (2001).

Three points can be taken from this discussion. First, although the overlap is not perfect, in all three countries examined here, our estimates typically match household survey-based estimates closely and are statistically indistinguishable. Second, we have noted a level of precision of survey-based estimates that is generally not terribly high. Third, standard errors on our estimators at the stratum-level are uniformly lower than those obtained with household survey data alone. This implies that errors introduced by applying the statistical

procedure outlined above are more than offset by the removal of sampling error when producing poverty estimates in the population census. We shall show next that it is possible to produce estimates of poverty with census data at levels of disaggregation far below what is possible with household survey data without paying any additional price in terms of statistical precision.

Table 3
Stratum-level poverty rates in South Africa (headcount)

Stratum	Household Survey(s.e.)	Census (s.e)
Western Cape	0.12 (0.011)	0.11 (0.006)
Eastern Cape	0.45 (0.014)	0.40 (0.009)
Northern Cape	0.38 (0.030)	0.35 (0.014)
Free State	0.51 (0.022)	0.53 (0.010)
Kwazulu-Natal	0.24 (0.014)	0.25 (0.008)
Northwest Province	0.37 (0.024)	0.41 (0.011)
Gauteng	0.11 (0.012)	0.17 (0.008)
Mpumalanga	0.26 (0.022)	0.22 (0.011)
Northern Province	0.36 (0.021)	0.35 (0.015)

Source: IES/OHS (1995); South Africa Census (1996).

Note: Standard errors on poverty estimates from the household survey reflect a 2-stage sampling design effect. Standard errors on poverty estimates from the census have been calculated according to the procedure outlined in the text.

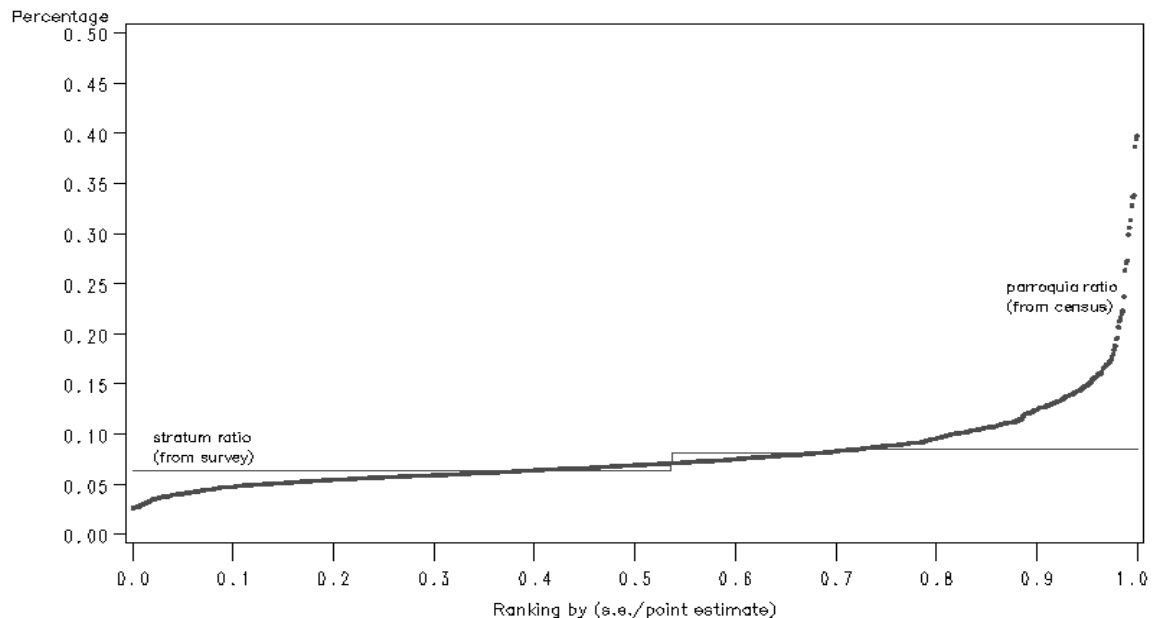
5.2 How low can we go?

The question of how far one can disaggregate in the population census depends on what is judged to be an acceptable level of statistical precision. As described above, and explored in greater detail in Elbers, Lanjouw, and Lanjouw (2001), the idiosyncratic component of the error in our estimator increases proportionately as the number of households in the target population falls. Any attempt to identify poor households in the census, for example, would be ill-advised because confidence bounds on household level poverty estimates would likely encompass the entire range between 0 and 1. However, the idiosyncratic error declines markedly as one aggregates across households, such that overall standard errors quickly become quite reasonable when estimates are made at the level of towns or districts. In Figures 1-9 below, it is shown that if one takes as a benchmark the precision which is achieved with household survey data at the representative stratum-level, then in all three

countries examined here, it is possible to produce estimates of poverty at the third administrative level (corresponding to 1000-2000 households on average in Ecuador and Madagascar, and 20,000 or so in South Africa) with similar levels of precision.

Figure 1

Standard Error as Percentage of Point Estimate: Rural Ecuador, Headcount
 Comparing survey based stratum-level estimates to census based parroquia-level estimates
 (915 Parroquias with an average 1050 households)



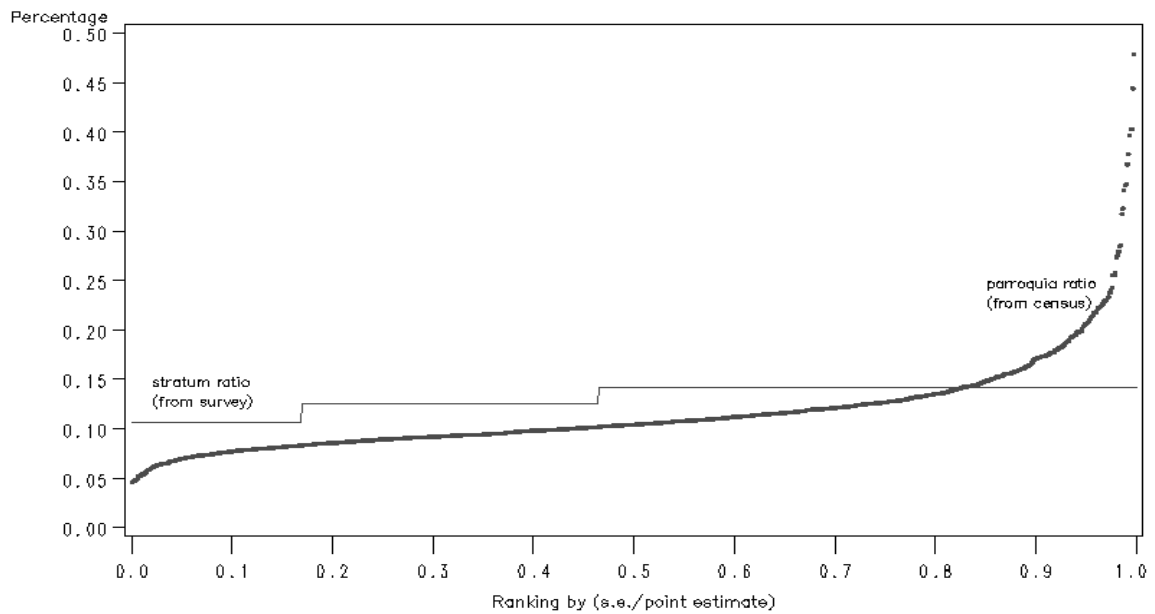
Source: authors' calculations.

Figure 1 illustrates the case for the headcount in rural Ecuador. We calculate the ratio of the standard-error to the point estimate for each of the 915 *parroquias* in rural Ecuador. The value of this ratio is represented by the vertical axis, and *parroquias* are ranked from lowest to highest along the horizontal axis. We overlay in this graph the value of the ratio from the survey estimates for the three strata covering rural Ecuador.¹⁶ From this figure we can see that for nearly 80 percent of *parroquias* the standard error as a share of the parroquia level poverty estimate is no higher than that typically found in household surveys. If we take the survey based stratum-level precision as a benchmark, such that the zone of ‘acceptability’ is up to the highest ratio value from survey estimates, we find that estimating poverty at this level of disaggregation does not result in particularly noisy estimates for the large majority of *parroquias* in the country. It should also be noted that in those cases for which the ratio is well above the survey level threshold, this usually occurs for those *parroquias* with particularly low poverty rates (standard errors decline as estimates decline, but not as sharply).

¹⁶ We compare ratios rather than absolute standard errors because we want to abstract away from the much greater variation in poverty estimated at the parroquia level compared to estimates at the stratum level from the household survey. Parroquias with very high estimated poverty tend to have larger standard errors, and the converse is the case for those with low poverty.

Figure 2

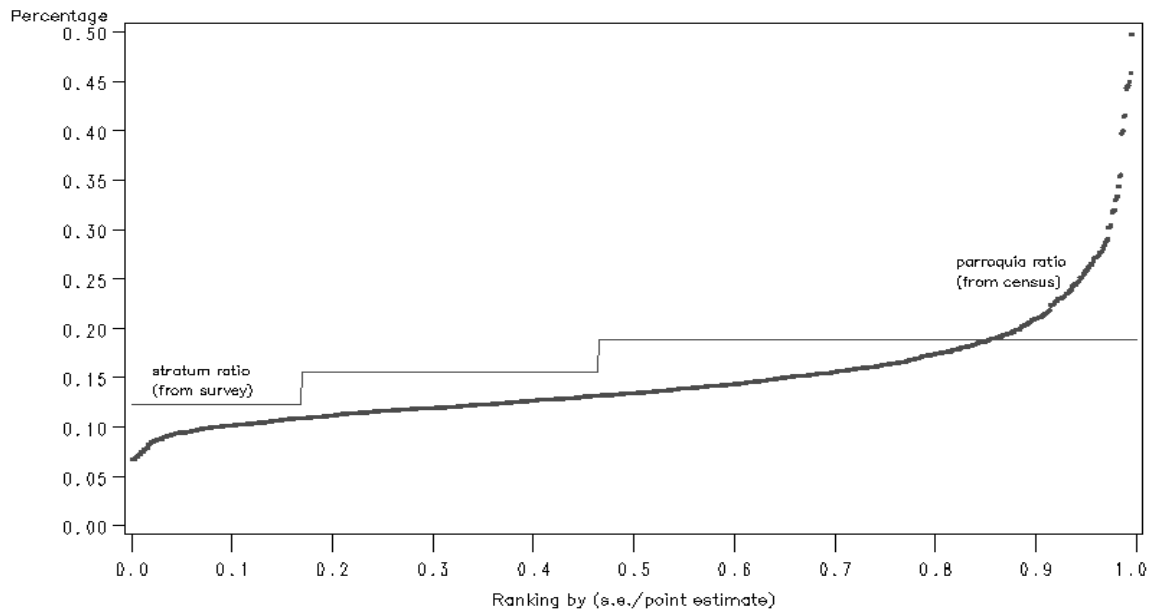
Standard Error as Percentage of Point Estimate: Rural Ecuador, FGT1
Comparing survey based stratum-level estimates to census based parroquia-level estimates
(915 Parroquias with an average 1050 households)



Source: authors' calculations.

Figure 3

Standard Error as Percentage of Point Estimate: Rural Ecuador, FGT2
Comparing survey based stratum-level estimates to census based parroquia-level estimates
(915 Parroquias with an average 1050 households)



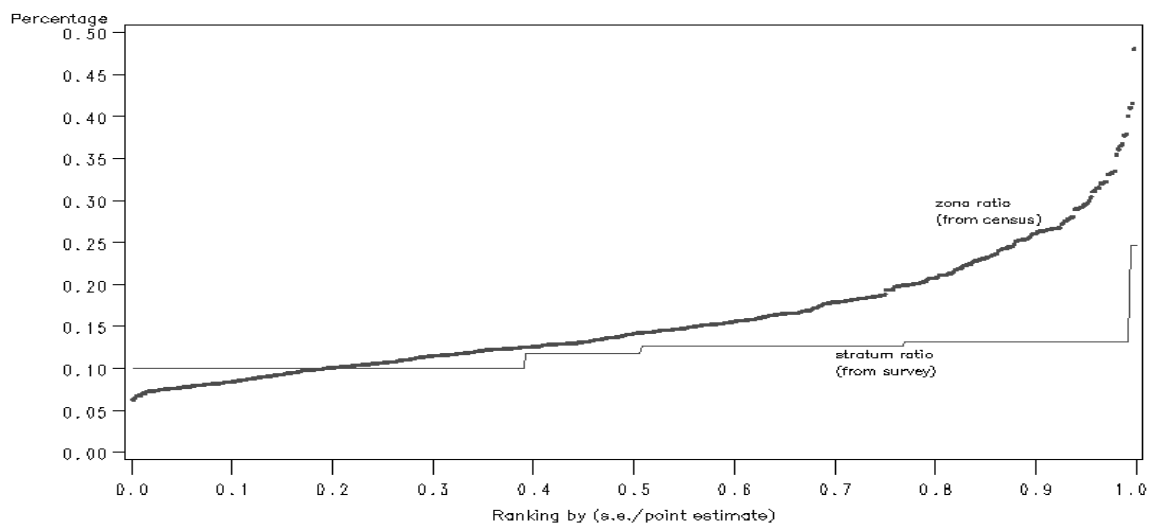
Source: authors' calculations.

Figures 2 and 3 show that the percentage of parroquias with lower 'standard error shares' than observed from the household survey increases when we consider poverty measures with greater distributional sensitivity than the headcount (FGT1 and FGT2). Comparing Figure 1 to Figures 2 and 3 we can see that this is not because FGT1 and FGT2 measures

are estimated with greater precision in the census, but rather survey-based estimates of these measures are less precise than of the headcount. It is remarkable that, for the FGT2 for example, nearly 90 percent of *parroquia* estimates of poverty are more precise than the corresponding stratum-level estimates of the FGT2 in the sample survey.

Figure 4

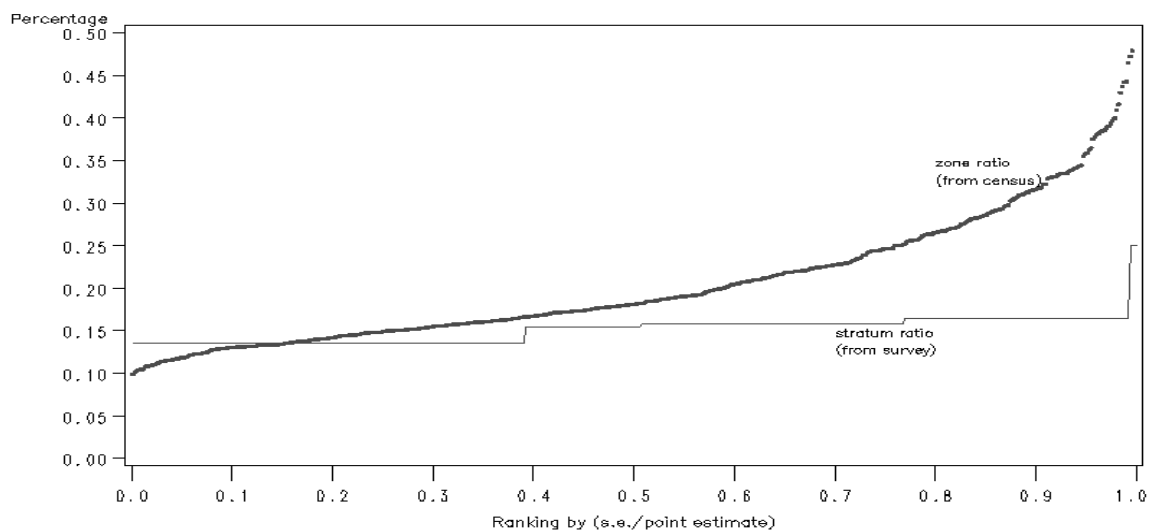
Standard Error as Percentage of Point Estimate: Urban Ecuador, Headcount
 Comparing survey based stratum-level estimates to census based zona-level estimates
 (453 Zonas with an average 1360 households)



Source: authors' calculations.

Figure 5

Standard Error as Percentage of Point Estimate: Urban Ecuador, FGT1
 Comparing survey based stratum-level estimates to census based zona-level estimates
 (453 Zonas with an average 1360 households)

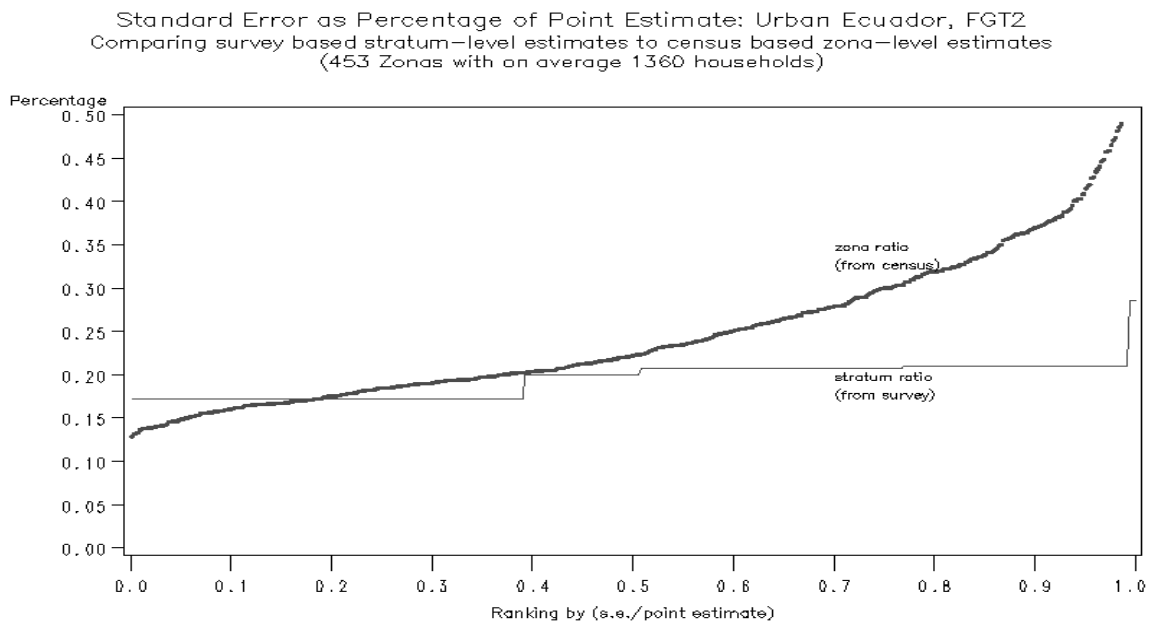


Source: authors' calculations.

In urban Ecuador the lowest administrative level is the 'zona' (roughly a neighbourhood). With the exception of one stratum (the Urban Oriente stratum—see Table 1) survey level standard error ratios are lower than for most zonas (Figures 4-6). Despite their small size, however, for a fairly large number of zonas the standard error ratios are not much higher

than their stratum-level counterparts in the survey. If one were to find these zona level ratios excessive the appropriate response would be to raise the level of aggregation when estimating urban poverty rates. Neighboring zonas could be joined into slightly larger groupings. While the picture is somewhat better with higher order FGT measures, the message remains that the zona is probably too low a level of disaggregation for urban areas in Ecuador.

Figure 6



Source: authors' calculations.

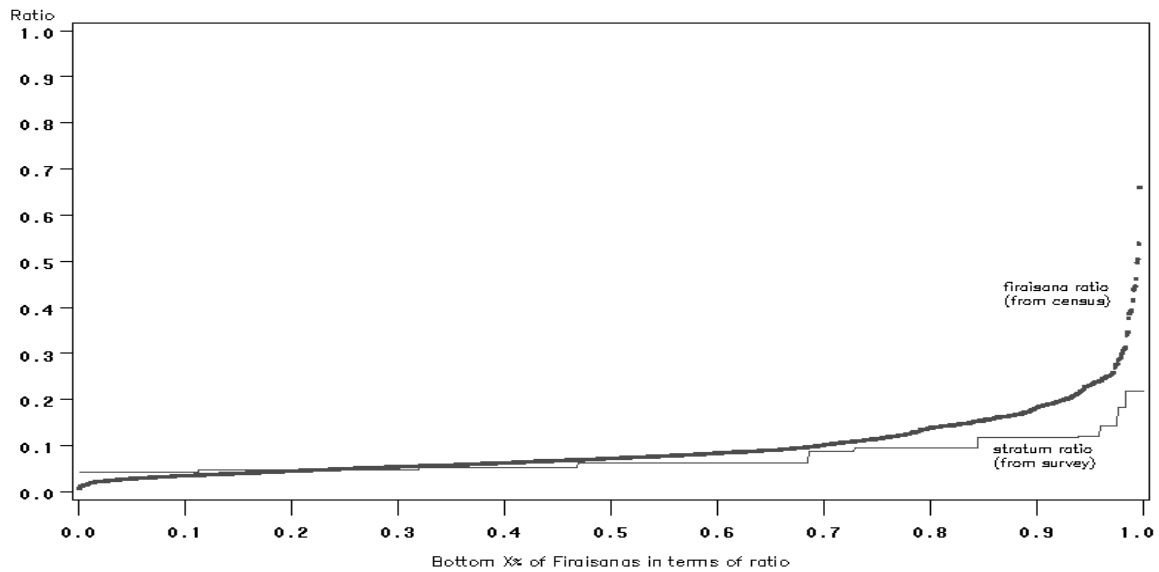
In Figure 7, we reproduce for Madagascar a similar picture as in Figure 1. The Firaisana is now the level of disaggregation (average number of households: 2000). Given that the sample estimates have ratios of standard error to point estimate as high as 20 percent (see Table 2) the vast majority of Firaisana level estimates look at least as good. Once again, if analysts are satisfied with the stratum-level precision obtainable with the EPM survey in Madagascar, then there should be no concern in working with Firaisana level estimates from the census.

In Figure 8 we see that the situation in South Africa is somewhat different. We start by disaggregating to the police-station level (with an average of 7,500 households).¹⁷ Here, the ratio of standard error to point estimates for police stations uniformly lies well above the ratio that obtains with the household survey at the stratum-level. Going down to police stations would require that an analyst be prepared to pay a price in terms of statistical precision of poverty estimates. When the level of disaggregation is to, say, the Magisterial District level (of which there are 354 in South Africa, with an average of 20,000

¹⁷ The police station does not correspond to a government administrative level, but comes closest to the third administrative level identified in Ecuador and Madagascar, in terms of population size.

Figure 7

Ratio of Estimated Standard Error to Point Estimate: Headcount
Firaïsona-level estimates in Madagascar
(1248 Firaïsonas; average number of households per Firaïsona: 1950)

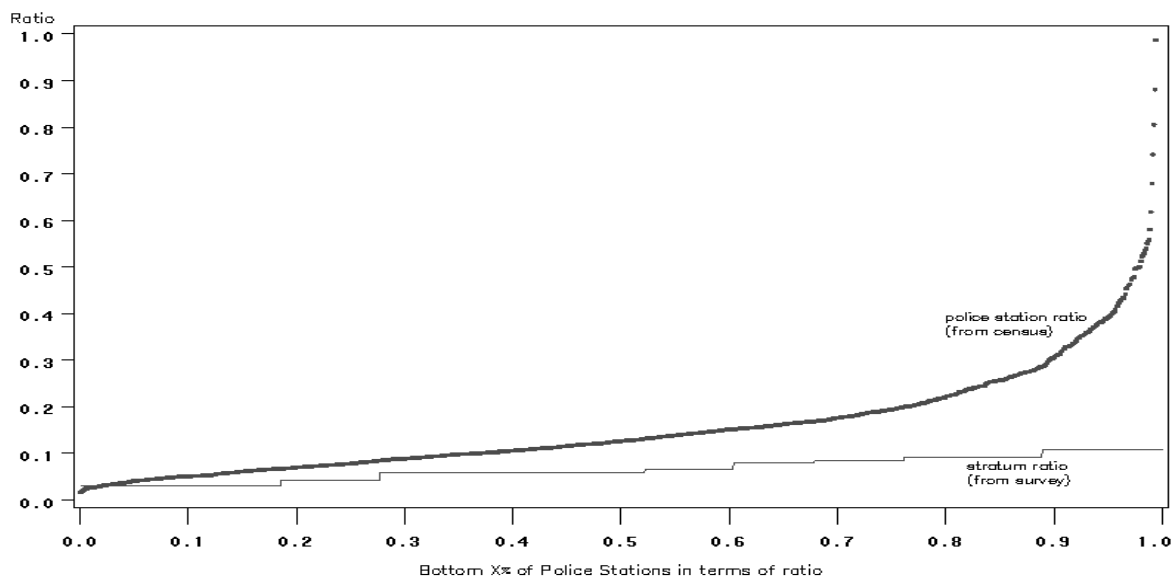


Source: authors' calculations.

households in each) the price is modest (see Figure 9). It is important to note, however, that the stratum-level estimates available with the South African IES survey are remarkably precise, because of the survey's large sample size (nearly 30,000 households). If one were to apply to the South African case the same standards of acceptability that are usually applied to settings where LSMS-style surveys prevail, even police station estimates of poverty would be viewed as remarkably precise.

Figure 8

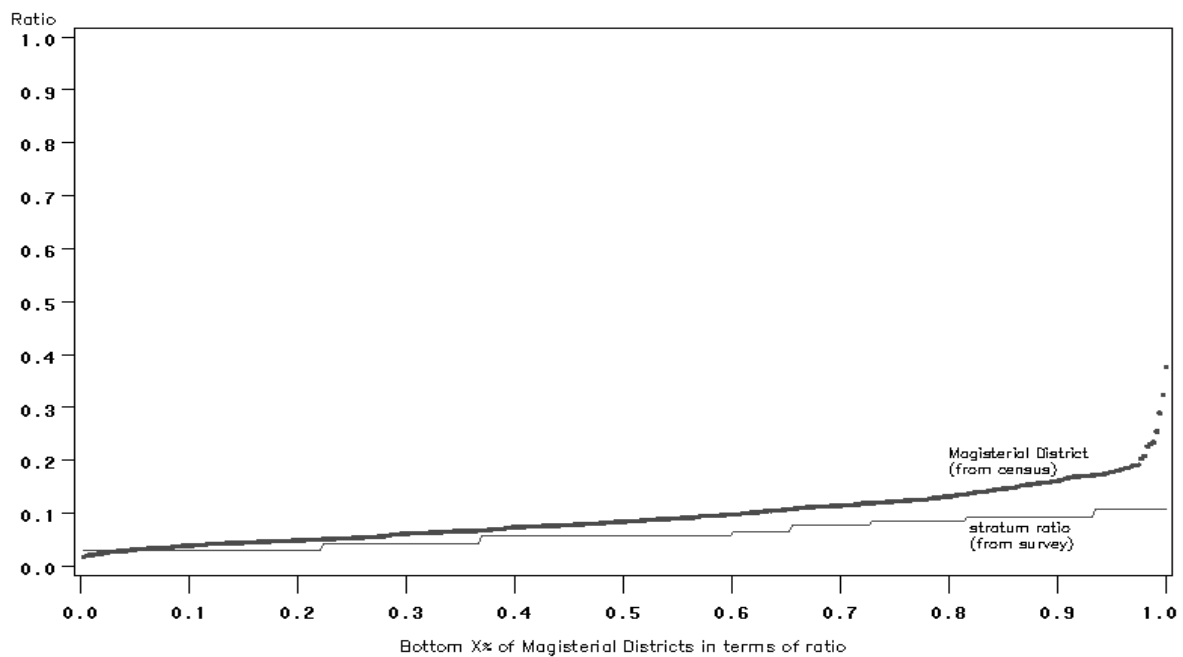
Ratio of Estimated Standard Error to Point Estimate: Headcount
Police station-level estimates in South Africa
(1096 Police Stations; average number of households per police station: 7500)



Source: authors' calculations.

Figure 9

Ratio of Estimated Standard Error to Point Estimate: Headcount
Magisterial District-level estimates in South Africa
(354 Magisterial Districts; average number of households per MD: 24000)



Source: authors' calculations.

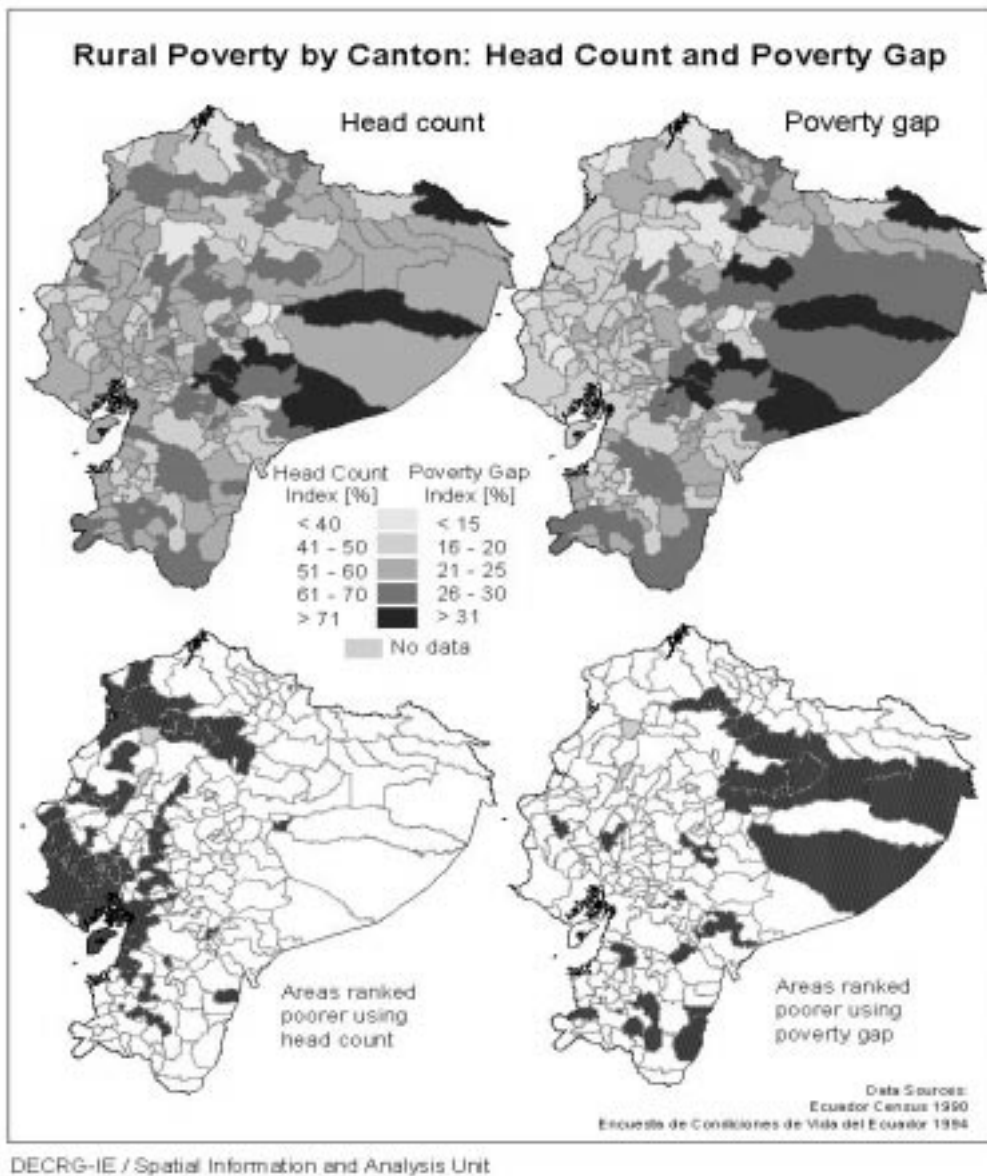
5.3 Geographic profiles: poverty in Ecuador, Madagascar and South Africa

The previous subsection has shown that estimates of poverty can be produced in our three example countries at levels of spatial disaggregation representing groupings of 1,000-20,000 households. Clearly, intermediate levels of spatial disaggregation are also possible. The question often arises how best to present information on the spatial distribution of poverty in a country once the number of estimates is large. A convenient manner in which to present the geographic poverty profile is in the form of maps where shadings are used to depict different degrees of poverty. Recent advances in digitized geographic information systems (GIS) have greatly facilitated the process of producing maps and offers great opportunities to combine the spatially referenced poverty information with other similarly referenced data. We illustrate here with a few examples some of the ways in which the spatially disaggregated poverty estimates produced with this methodology can be represented in map form.

Figure 10 displays the spatial distribution of estimated rural poverty in Ecuador at the *cantonal* level (second administrative level representing around 5,000-10,000 households). Comparisons between the Costa, the coastal region of Ecuador, and the Sierra, the central mountainous region, feature highly in popular political debate in Ecuador.¹⁸ The top two maps in Figure 10 depict the spatial distribution of poverty on the basis of two common

¹⁸ See for example, 'Under the Volcano', *The Economist*, 27 November 1999:66.

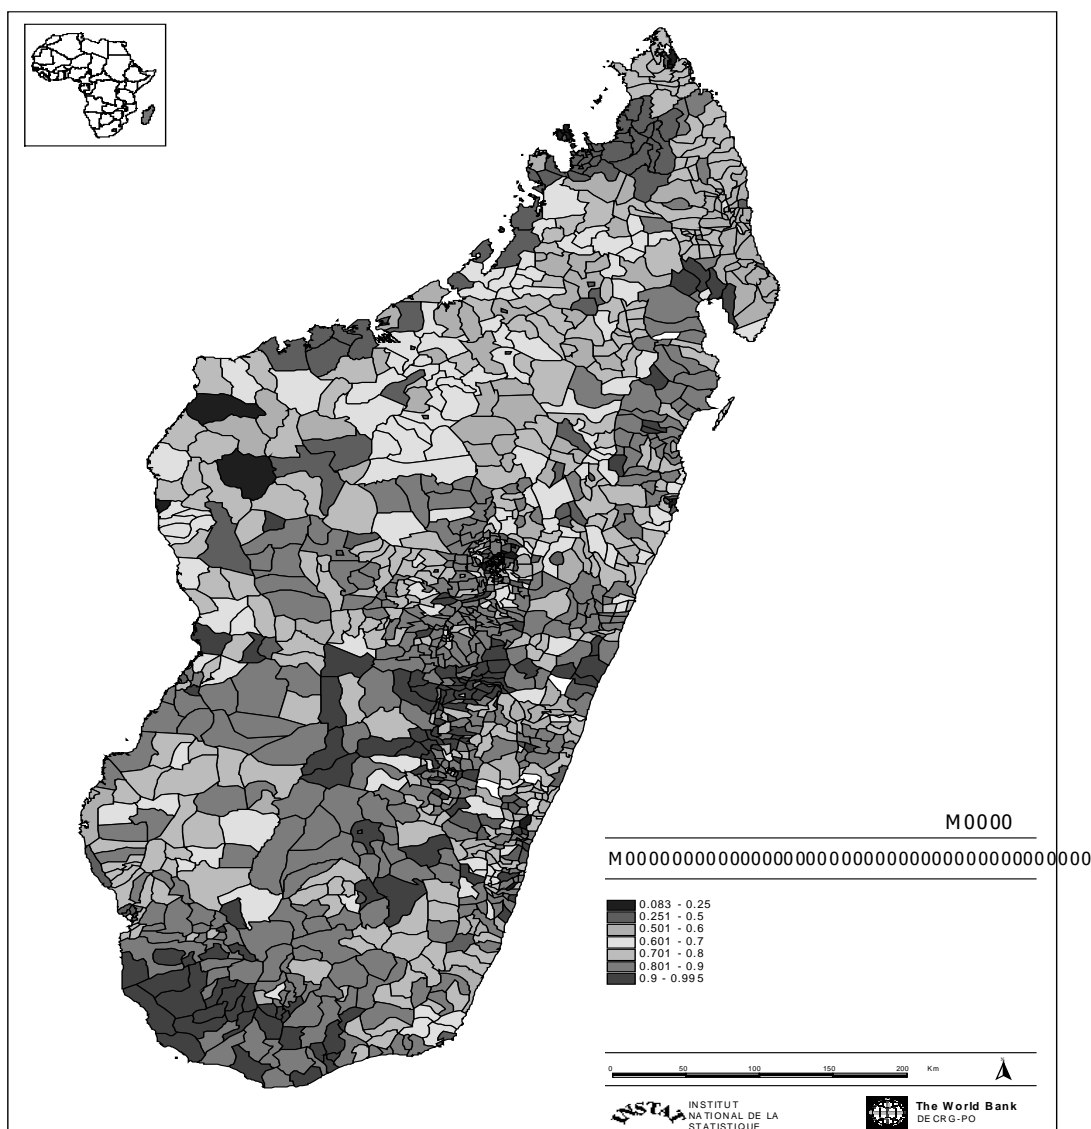
Figure 10



Source: authors' calculations.

measures: the headcount and the poverty gap (FGT1). The bottom two maps in Figure 10 indicate those instances where the two alternative poverty measures differ in their ranking of cantons. The map on the lower left shows that in the Costa a number of cantons are ranked poorer under the headcount criterion than under the poverty gap. In contrast, in the Sierra and the less populated east (Oriente), numerous cantons are ranked more poor under the poverty gap criterion than under the headcount. Clearly, views about the relative poverty of the regions will be affected by the measure of poverty employed. The discussion in this paper has placed considerable emphasis on statistical precision of poverty estimates produced with the methodology outlined here. As one thinks about drawing maps describing the spatial distribution of poverty, it is also important to convey information about statistical precision in those maps. Figures 11 and 12 are an attempt to do so for Madagascar. Figure 11 displays our geographic poverty profile for over 1300 firaisanas in

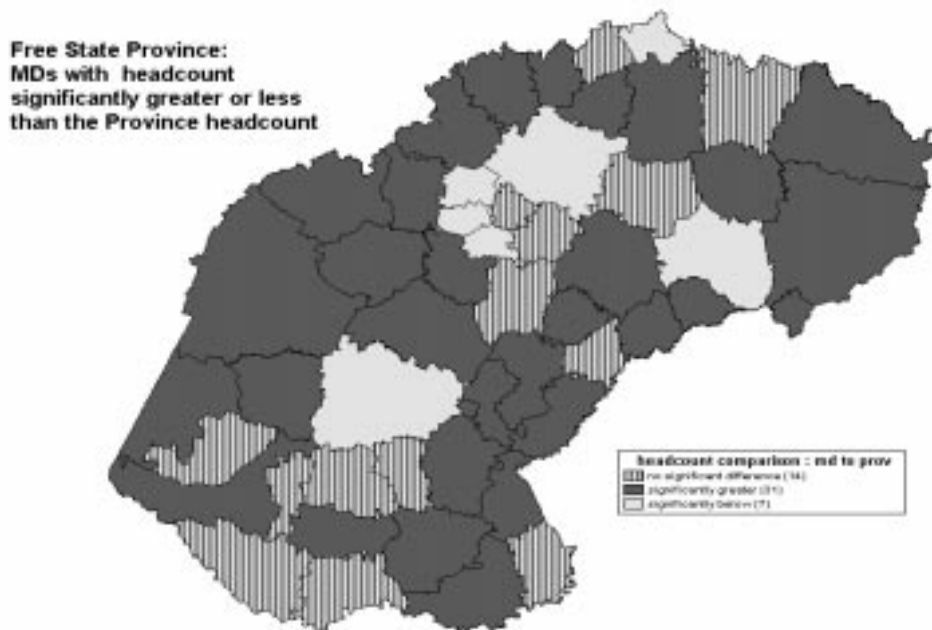
Figure 11
Firaisana Level FGT(0) Estimates



Source: authors' calculations.

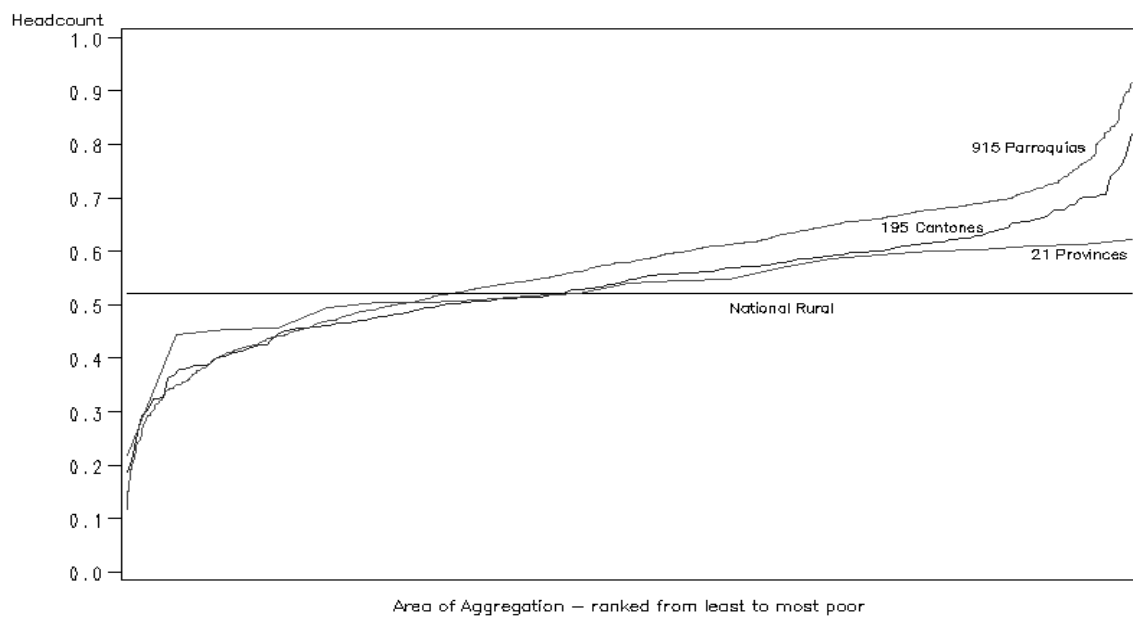
Madagascar and Figure 12 shows that 68 percent of the 1332 firaisanas in Madagascar have headcount rates that are significantly different than the headcount rate for the stratum to which they belong. Figure 13 indicates that within South Africa's poorest province, Free State Province, poverty is not homogeneously distributed. A number of MD's within this province record an incidence of poverty that is significantly lower than that of the province overall and others are considerably more poor. This observation follows directly from the fact that poverty measures such as the headcount, poverty gap (FGT1) and squared poverty gap (FGT2) all belong to a class of subgroup decomposable poverty measures (Foster, Greer and Thorbecke index, 1984). The poverty rate for a given locality is equal to the population weighted average of poverty rates of sublocalities located within that area. Because the poverty rate for the given locality is an average, it is clear that some

Figure 13
Poverty within poverty: South Africa



Source: authors' calculations.

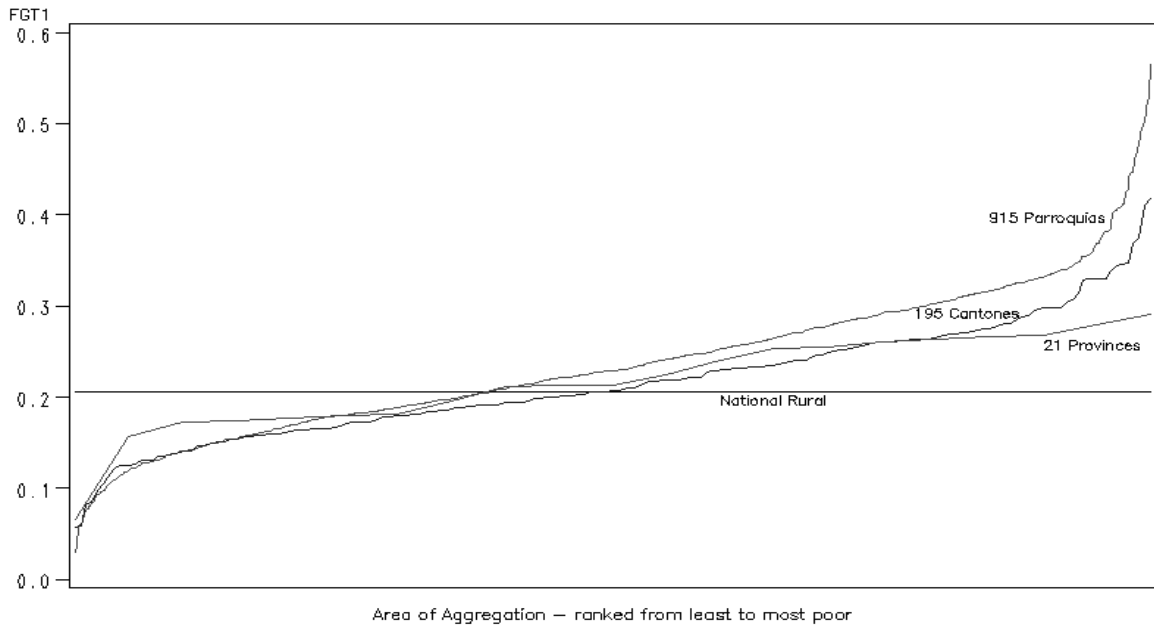
Figure 14
Poverty by Area of Aggregation: Headcount
Comparing Parroquia versus Cantonal versus Provincial-level estimates in rural Ecuador
Areas of aggregation ranked from least poor to most poor



Source: authors' calculations.

Figure 15

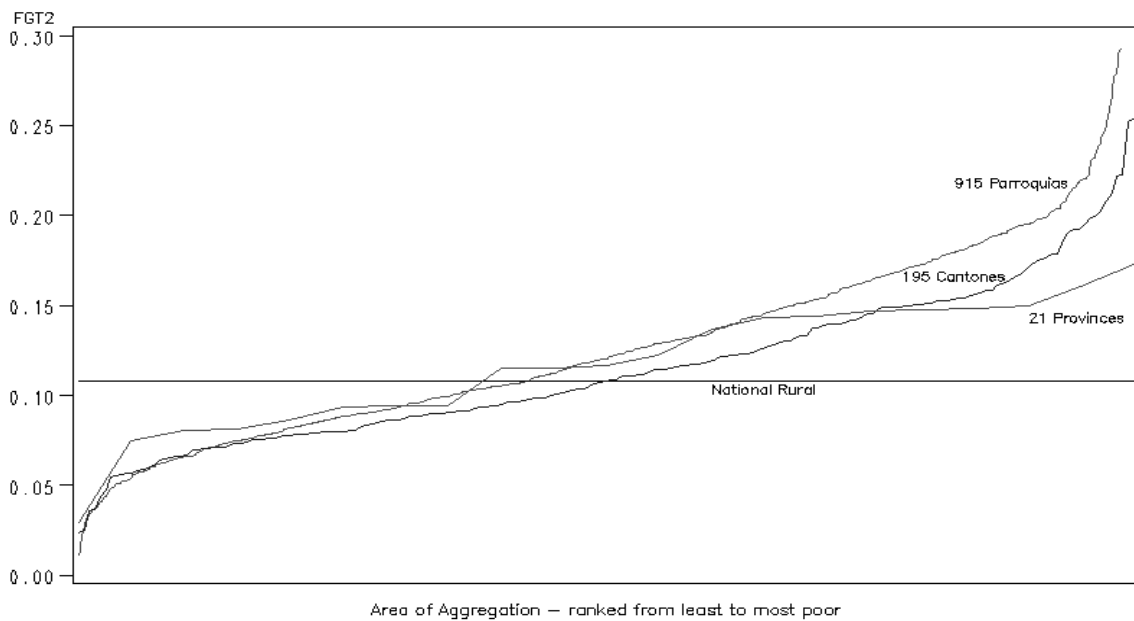
Poverty by Area of Aggregation: FGT1
Comparing Parroquia versus Cantonal versus Provincial-level estimates in rural Ecuador
Areas of aggregation ranked from least poor to most poor



Source: authors' calculations.

Figure 16

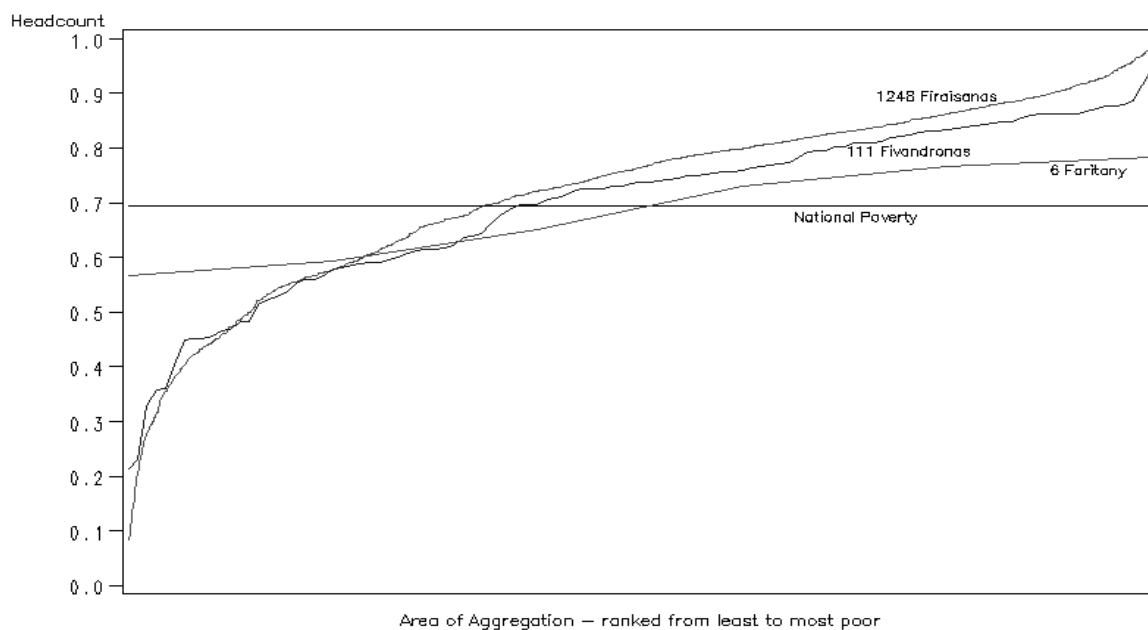
Poverty by Area of Aggregation: FGT2
Comparing Parroquia versus Cantonal versus Provincial-level estimates in rural Ecuador
Areas of aggregation ranked from least poor to most poor



Source: authors' calculations.

Figure 17

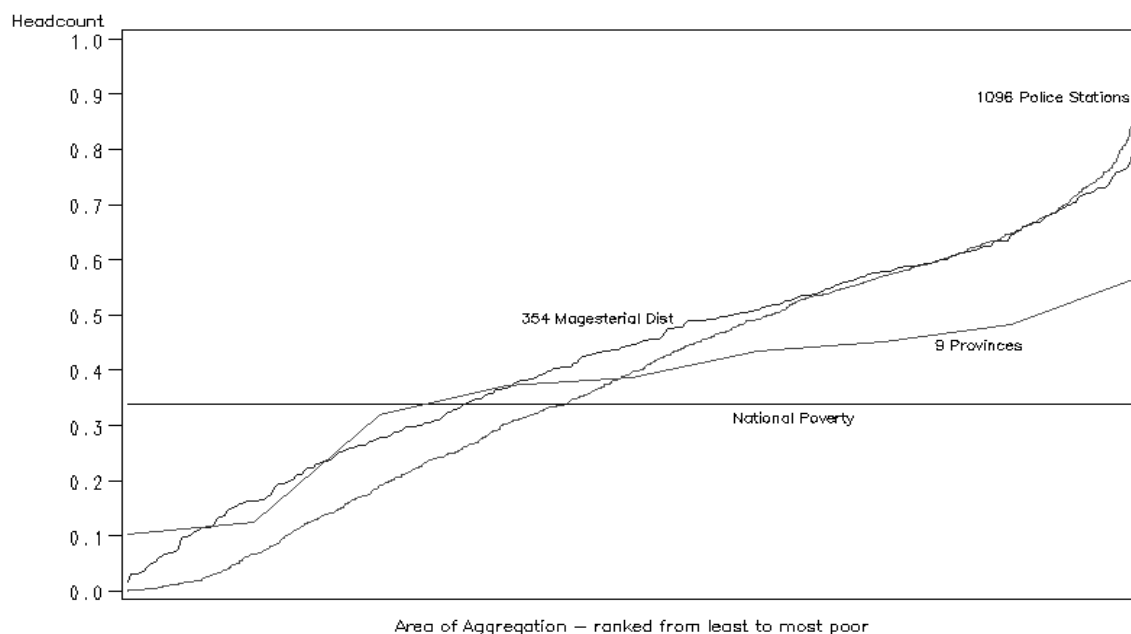
Poverty by Area of Aggregation: Headcount
Comparing Firaiana versus Fivandrona versus Faritany—level estimates in Madagascar
Areas of aggregation ranked from least poor to most poor



Source: authors' calculations.

Figure 18

Poverty by Area of Aggregation: Headcount
Comparing Police Station versus Magisterial District versus Province—level estimates in South Africa
Areas of aggregation ranked from least poor to most poor



Source: authors' calculations.

Figures 14-18 illustrate to what extent this observation holds in Ecuador, Madagascar and South Africa. Figure 14 ranks localities in rural Ecuador by incidence of poverty—in turn

provinces, cantons and parroquias—and examines the spread of poverty of localities around the national level. This spread is lowest for provinces, followed by cantons and then parroquias.¹⁹ The same pattern obtains for the FGT1 and FGT2 poverty measures (Figures 15 and 16, respectively).

In Madagascar while the pattern is again the same, it is noteworthy that the degree to which fivandrona level estimates understate the spatial variation in poverty relative to firaisana level estimates is not that great—despite there being more than ten times as many firaisanas as fivandronas (Figure 17). Similarly, the 354 Magesterial Districts in South Africa also do a fairly good job of capturing the variation in poverty that the 1096 police station level estimates depict (Figure 18).

6 Conclusions

This paper has taken three developing countries, Ecuador, Madagascar and South Africa, and has implemented in each a methodology to produce estimates of poverty at a level of disaggregation that to date has not generally been available. The countries are very unlike each other—with different geographies, stages of development, quality and types of data, and so on. Nonetheless the paper has demonstrated that the methodology works well in all three settings and can be seen to produce valuable information about the spatial distribution of poverty within those countries that was previously not known.

The methodology is based on a statistical procedure to combine household survey data with population census data, by imputing into the latter a measure of economic welfare (consumption expenditure in our examples) from the former. The poverty rates that are produced are *estimates* and as such are subject to statistical error. The paper has demonstrated that the poverty estimates produced from census data are plausible in that they match well the estimates calculated directly from the country's surveys (at levels of disaggregation that the survey can bear). The precision of the poverty estimates produced with this methodology vary with the degree of disaggregation. We have shown that in all three countries considered here our poverty estimators allow one to work at a level of disaggregation far below that allowed by surveys.

We have illustrated how the poverty estimates produced with this method can be represented in maps, thereby conveying an enormous amount of information about the spread and relative magnitude of poverty across localities, as well as the precision of estimates, in a way which is quickly and intuitively absorbed also by non-technical audiences. Such detailed geographical profiles of poverty can inform a wide variety of debates and deliberations, amongst policymakers as well as civil society.

¹⁹ Note that as we are working with *expected* poverty rather than actual measures, the *true* spread of poverty, for any given level of disaggregation, is likely to be larger than that which we observe.

We have finally noted that perceptions as to the importance of geographical dimensions of poverty are themselves a function of the degree of spatial disaggregation of available estimates of poverty. The smaller the localities into which a country can be broken down the more likely one will conclude that geography matters.

References

- Aghion, P. and Bolton, P. (1997) 'A Theory of Trickle-Down Growth and Development', *Review of Economic Studies* 64(2):151-72.
- Alderman, H., Babita, M., Demombynes, G., Makhatha, N. and Özler, B. (2001) 'How Low Can You Go? Combining Census and Survey Data for Mapping Poverty in South Africa', mimeo, DECRG-World Bank: Washington DC.
- Alesina, A. and Rodrik, D. (1994) 'Distributive Politics and Economic Growth', *Quarterly Journal of Economics* 109:465-90.
- Banerjee, A. and Newman, A. (1993) 'Occupational Choice and the Process of Development', *Journal of Political Economy* 101(1):274-98.
- Deaton, A. (1999) 'Inequalities in Income and in Health', *NBER Working Papers* No.7141.
- Demombynes, G. and Özler, B. (2001) 'Inequality and Crime in South Africa?', mimeo, DECRG-World Bank: Washington DC.
- Elbers, C., Lanjouw, J.O. and Lanjouw, P. (2001) 'Welfare in Villages and Towns: Micro level Estimation of Poverty and Inequality', mimeo, DECRG-World Bank: Washington DC.
- Elbers, C., Lanjouw, J.O., Lanjouw, P. and Leite, P.G. (2001) 'Poverty and Inequality in Brazil: New Estimates from Combined PPV-PNAD Data', mimeo, DECRG-World Bank: Washington DC.
- Galor, O. and Zeira, J. (1993) 'Income Distribution and Macroeconomics', *Review of Economic Studies* 60:35-52.
- Hentschel, J. and Lanjouw, P. (1996) 'Constructing an Indicator of Consumption for the Analysis of Poverty: Principles and Illustrations with Reference to Ecuador', *LSMS Working Paper* No.124, DECRG-World Bank: Washington DC.
- Hentschel, J., Lanjouw, J.O., Lanjouw, P., and Poggi, J. (2000) 'Combining Census and Survey Data to Trace the Spatial Dimensions of Poverty: A Case Study of Ecuador', *World Bank Economic Review* 14(1)147-65.
- Mistiaen, J., Özler, B., Razafimanantena, T., and Razafindravonona, J. (2001) 'Disaggregated Maps of Estimated Poverty and Inequality for Madagascar in 1993'. mimeo, DECRG-World Bank: Washington DC.

- Murphy, K.M., Schleifer, A., and Vishny, R. (1989) 'Income Distribution, Market Size and Industrialization', *Quarterly Journal of Economics* 104:537-64.
- Persson, T. and Tabellini, G. (1994) 'Is Inequality Harmful for Growth', *American Economic Review* 84:600-21.
- World Bank (1996) 'Ecuador Poverty Report', *World Bank Country Study*, Ecuador Country Department, World Bank: Washington DC.