The Geography of Poverty in Morocco:
Micro-Level Estimates of Poverty and Inequality from Combined Census and Household Survey Data

Peter Lanjouw (DECRG)
February 24, 2004

Introduction

Welfare levels tend to vary among the regions of almost every country of the world. Morocco is no exception: pockets of severe deprivation are a widely acknowledged, albeit only partially documented, phenomenon.[1] The existence of such poor areas can be due to differences in geographic capital-biophysical endowment, access to infrastructure and markets, etc.-as well as due to government policies, such as the distribution of centrally allocated resources, or migration policies.

In the face of such geographic heterogeneity, successful policy making in Morocco requires a good information base. For instance, an understanding of poverty and inequality levels at detailed spatial scales is a prerequisite for fine geographic targeting of interventions aimed at improving welfare levels. Decentralization has meant that decision making for poverty alleviation programs is shifting from central government to regional or local levels. However, local decision making, the design of the decentralization processes and even the decision whether or not to pursue further decentralization, should be based on reliable, locally relevant information on living standards and the distribution of wealth. In Morocco, to date, such information is not readily to hand.

The problem of a lack of locally relevant poverty information is common. There are two main types of welfare-related information sources available to policy-makers. *Household surveys* often include a detailed income and/or consumption expenditure module. However, due to relatively small sample size, the collected information is usually only

[1] See for example, Repères Statistiques, Bulletin Mensuel de la Direction de la Statistique, No. 77, Janvier, 2003, Département de la Prévesion Economique et du Plan, Rabat.

representative for broad regions of the country. *Census data* (and sometimes large household sample surveys) are available for all households (or very large samples of households) and can provide reliable estimates at highly disaggregated levels such as small municipalities, towns, and villages. But censuses do not contain the necessary income or consumption information to yield reliable indicators of the level and distribution of welfare such as poverty rates or inequality measures.

Recent research has explored a technique that addresses the problem of lack of local data on poverty and inequality.[2] This method combines survey and census data to estimate consumption-based welfare indicators for small geographic areas such as provinces and communes-i.e., a *poverty map*. The method has been implemented in a growing number of developing countries and experience from these efforts suggests that statistically reliable estimates of poverty and inequality are attainable at encouragingly fine levels of spatial detail.

Methodology: An Outline of the Basic Steps to Poverty Mapping

The poverty map methodology involves detailed analysis of two main sources of data: a household survey; and the population census. In the first stage of the analysis the two data sources are subjected to very close scrutiny with an eye towards identifying a set of common variables. In the second stage the survey is used to develop a series of statistical models which relate per capita consumption to the set of common variables identified in the preceding step. In the final stage of the analysis the parameter estimates from the previous stage are applied to the population census and used to predict consumption for each household in the population census. Once such a predicted consumption measure is available for each household in the census, summary measures of poverty (and/or inequality) can be estimated for a set of households in the census. Statistical tests can be performed to assess the reliability of the poverty estimates that have been produced.

[2] Elbers, Lanjouw and Lanjouw, 2002 and 2003, refine and extend considerably an approach first outlined in Hentschel, Lanjouw, Lanjouw and Poggi, 2000.

The three stages of analysis occur in sequence. The *first stage* in the poverty mapping exercise involves a rather painstaking comparison of common variables across the household survey and the population census. The idea here is to identify variables at the household level that are defined in the same way in both the household survey and the census. It is important that this common subset of variables be defined in exactly the same way across the two data sources. One can check for this rigorously on the basis of statistical tests of differences, or informally by simply making visual comparisons.

A concurrent exercise that can be carried out in parallel to the exercise described above is the compilation of a database at a level of aggregation higher than the household, which can be inserted into the household level census and the household survey databases. One of the methodological concerns in the poverty mapping exercise is that the common pool of household variables will not suffice to capture unobserved geographic effects, such as agro-climatic conditions, quality of local government administration, etc., which might still be very important in predicting household level consumption. In an effort to proxy these unobserved factors, it will be useful to merge, say, district or sub-district data that have been compiled separately into both the census and the household survey. These sub-district level data may comprise a wide range of variables (for example, construction of schools, public spending figures, infrastructure availability, population estimates, etc.)

If the two tasks described above yield a good and reasonably large set of common household-level variables, supplemented by a series of additional variables at a slightly higher level of disaggregation, the decision can then be made to proceed to the *second-stage* analysis. This stage involves the econometric estimation of models of consumption on the set of household-level and community variables. Different models are estimated for each stratum in the household survey data set, separately for rural and urban areas.

Successful completion of the *second-stage* analysis permits one to take the parameter estimates and attendant statistical outputs to the *third stage*. This stage is associated with the imputation of consumption into the census data at the household level, and the

estimation of poverty and inequality measures at a variety of levels of spatial disaggregation. Statistical precision of the welfare estimates is also gauged in this stage.

Once the poverty map exercise has been completed for all regions in the country, the resulting databases which provide estimates of poverty and inequality (and their standard errors) at a variety of levels of geographic disaggregation can be projected onto geographic *maps* using GIS (Geographic Information Systems) mapping techniques. This involves the application of GIS software (such as ARCView) which merges information on the geographic coordinates of localities such as the district or sub-district with the poverty and inequality estimates produced by the poverty mapping methodology.

Mapping Poverty in Morocco

The recently completed Moroccan poverty map has been produced on the basis of the methodology outlined above. Certain details of the Moroccan context warrant specific mention.

Sources of Data: The poverty map in Morocco has been developed on the basis of the 1994 population census and the 1998 EPM household survey. Given that these two data sources do not refer to the same time period, interpretation of results is somewhat problematic: does the poverty map produced on the basis of these two datasets refer to 1994 or to 1998? This question is difficult to answer precisely. If there have been no major changes in the structure of the Moroccan economy between 1994 and 1998, such that returns to explanatory variables in the consumption models (factors like education, occupation, and so on) have not changed, then the map can be viewed as representing the spatial distribution of poverty in 1994 – the year for which levels of schooling, and other explanatory variables in the regression models, apply. If there have been major changes in returns to factors during this time period, then the map does not correspond clearly to either 1994 or 1998 but rather to some combination of these years. Clearly, interpretation is more straightforward if the former assumption, of little change in returns, holds. Of course, this assumption becomes quite plausible if the period between 1994 and 1998 is one of general stagnation, when for example, even levels of schooling, population,

occupational breakdown have remained constant. In this case the map refers not only to 1994 but also to 1998, and the depiction of the spatial distribution of poverty presented in this study can be viewed as valid over a longer period of time. In reality, changes in both levels of factors as well as their returns are likely to have taken place during the 1990s in Morocco. But these changes are not thought to have been dramatic. As a result, it does not seem unreasonable to view the poverty maps produced here as providing a snapshot of the geographic distribution of poverty for the period of the "mid to late 1990s". Additional methodological work is warranted to pursue these questions further.

Sampling structure of Moroccan Census. One of the features of the Moroccon population census data is that only 25% of census questionnaires were computerized. This sample was not selected on a purely random basis. Rather, provinces with small populations were entered completely, in medium sized provinces one in two questionnaires were entered, in larger provinces, one in four questionnaires were entered, and in the very large provinces only one in ten questionnaires was computerized. The implication of this feature of the Moroccan census data is that poverty map estimates produced for Morocco are likely to be slightly less precisely estimated than they would have been otherwise. However, because the population-related structure underpinning the census sample ensures that the smallest provinces are well represented in the data, the overall impact of this issue has been found to be rather modest.

Dropping all households "without residence". A decision was taken to drop from the census dataset all households that were not residing in a fixed residence ("sans abri"). This is due to the fact that the EPM household survey did not include such households, and as such the models estimated in the EPM are not likely to apply readily to these households. Moreover, a number of the explanatory variables in the EPM models are linked to characteristics of the house (access to water, etc.) and these are missing for the homeless segment of the population. Note that this segment of the population is not terribly large, as a rule, but in the very south of the country (in both urban and rural areas) it is somewhat more important quantitatively, presumably because of the existence of a nomadic population in this region. The population over which the poverty map estimates

presented in this report are believed to hold thus refers to a (large) subset of the overall population of Morocco.

Estimating Poverty in Regions not Covered by the EPM Survey. The Moroccan EPM household survey sample has been stratified down to the regional level, and distinguishes as well between urban and rural areas within each region. There are about 16 regions in Morocco, and given the limited sample size of the household survey (about 5000 households) this implies that there are some regions in which only few households were sampled (less than 100). Consumption models have thus been estimated in nine "domains" which are built on the basis of 3-5 geographically contiguous regions, separately for rural and urban areas. The southernmost geographic domain comprises regions 1, 2, 3, 4 and 7 (The regions of Oued Ed-Dahab-Lagouira, Layoune-Boujdour-Sakia El Hamra, Guelmin-Es-Semara, Souss-Massa-Draa, and Marrakech-Tensift-Al-Haouz, respectively). Regions 1 and 2 are in large desert regions with fairly sparse populations. The 1998 ENV survey did not sample in the first two regions and even region 3 comprised only a small sample of less than 100 households in total. Consumption models for this geographic grouping were thus based on regions 3, 4 and 7 only. We have imputed consumption to census data and estimated poverty rates for these two regions implicitly assuming that the model parameters that hold for regions 3, 4 and 7 also apply to regions 1 and 2. This is not an assumption we have been able to independently validate. Indeed, it is probably most sensible to treat the poverty map estimates for regions 1 and 2 separately from others, and regard them as considerably more tentative than the estimates produced for the rest of the country.

Poverty Lines Two poverty lines that have been employed in this exercise. The poverty line for rural areas comprises 3037 dirhams per person per year while that for urban areas corresponds to 3922 dirhams (in 1998 currency). Note, even with this considerably higher poverty line for urban areas, poverty in rural Morocco is generally much higher than in urban areas. These two poverty lines are best viewed as conservative, and are associated with a national poverty headcount rate of around 17%. To probe robustness of the geographic profile of poverty to alternative poverty lines, we have re-estimated

poverty measures (headcount, poverty gap and squared poverty gap) based on alternative poverty lines 50% higher than the benchmark lines mentioned above.

In addition to producing poverty estimates the Moroccan poverty mapping project has also estimated two additional measures of relevance to the analysis of welfare: average per capita consumption and per capita consumption inequality. The inequality measure that has been used for this purpose is the General Entropy class inequality measure with parameter value 0.5. This inequality measure is less computationally intensive to calculate than the more commonly used Gini coefficient, but has been found to be quite satisfactory for purposes of summarizing inequality.

Region, Province and Commune-level Estimates of Poverty in Morocco

How well do the poverty estimates, calculated from consumption data imputed at the household level into the population census, tally with poverty estimates calculated directly from the EPM household survey? The question cannot be answered at the level of disaggregation to which it will be desirable to produce estimates with census data. The household survey is, at best, representative only at the level of regions and cannot yield reliable estimates at the commune, or even province, level. Table 1 compares estimates deriving from household survey data versus those from census data derived with the poverty mapping methodology at the level of the geographic "domains" at which consumption models were estimated in the survey.

Table 1: Incidence of Poverty in Morocco:
Comparing Direct Survey Estimates Against Imputed Census-based Estimates

Domain of Analysis	EPM Survey Estimate	Estimated Poverty from Census	Population
Regions 3-4-7 Rural	0.227	0.232	3,574,876
Regions 3-4-7 Urban	0.075	0.089	1,998,199
Regions 5-13-14 Rural	0.289	0.277	2,260,372
Regions 5-13-14 Urban	0.181	0.221	2,457,006
Regions 6-9-10-11-12 Rural	0.169	0.179	3,529,200
Regions 6-10-11-12 Urban	0.069	0.094	3,163,858
Region 9 Urban	0.032	0.041	2,882,356
Region 8-15-16 Rural	0.213	0.244	2,940,842
Region 8-15-16 Urban	0.107	0.102	2,449,708
All-Morocco	0.153	0.166	25,256,417

Note:
Regions 3-4-7 refer to Guelmim-Es-Semara, Souss-Massa-Draa, and Marrakech-Tensift-Al-Haouz.
Regions 5-13-14 refer to Gharb-Chrarda-Beni-Hssen, Meknes-Tafilalet, and Fes-Boulemane
Regions 6-9-10-11-12 refer to Chaouia-Ouardigha, Grand-Casablanca, Rabat-Sale-Zemmour-Zaer, Doukkala-Abda, and Tadla-Azilal.
Region 9 refers to Grand Casblanca
Regions 8-15-16 refer to Oriental, Taza-Al Hoceima-Taounate, Tanger-Tetouan

Table 1 yields several important observations. First, at the level of "domains of analysis" estimates of poverty from the census match closely those from the household survey data. The Table does not report standard errors on the poverty estimates in the two columns, but incorporating this additional information yields the conclusion that in no case can one reject the hypothesis that the two poverty estimates for a given row entry are the same. Thus, the poverty mapping methodology appears to be yielding sensible estimates of poverty at this level of aggregation. Second, while not significantly different in a statistical sense, the models and parameter estimates produced in the poverty mapping exercise yield poverty estimates that are slightly higher than those obtained from the 1998 household survey alone. This is consistent with a slight improvement in general welfare levels over time. Third, the incidence of poverty in urban areas is consistently, and markedly, lower than in rural areas. This is a robust finding from survey-based analysis as well as from the poverty mapping exercise, and is generally repeated as one disaggregates further (see below).

The strength of the poverty mapping methodology lies, of course, not so much in terms of the estimates of poverty it yields at the level of "domains of analysis" or other similarly aggregated localities, but in terms of more disaggregated estimates of poverty. Maps 1 and 2 depict province-level poverty (rural and urban areas combined) and population density. Such maps are convenient means for transmitting a large amount of information in a fairly non-technical manner. An important insight from Map 1, for example, concerns the considerable degree of heterogeneity across provinces in estimated poverty: poverty rates between geographically contiguous provinces can vary markedly. There is evidence of both "pockets of poverty" and "islands of prosperity" at the province level.

Map 1: Incidence of Poverty at the Province Level

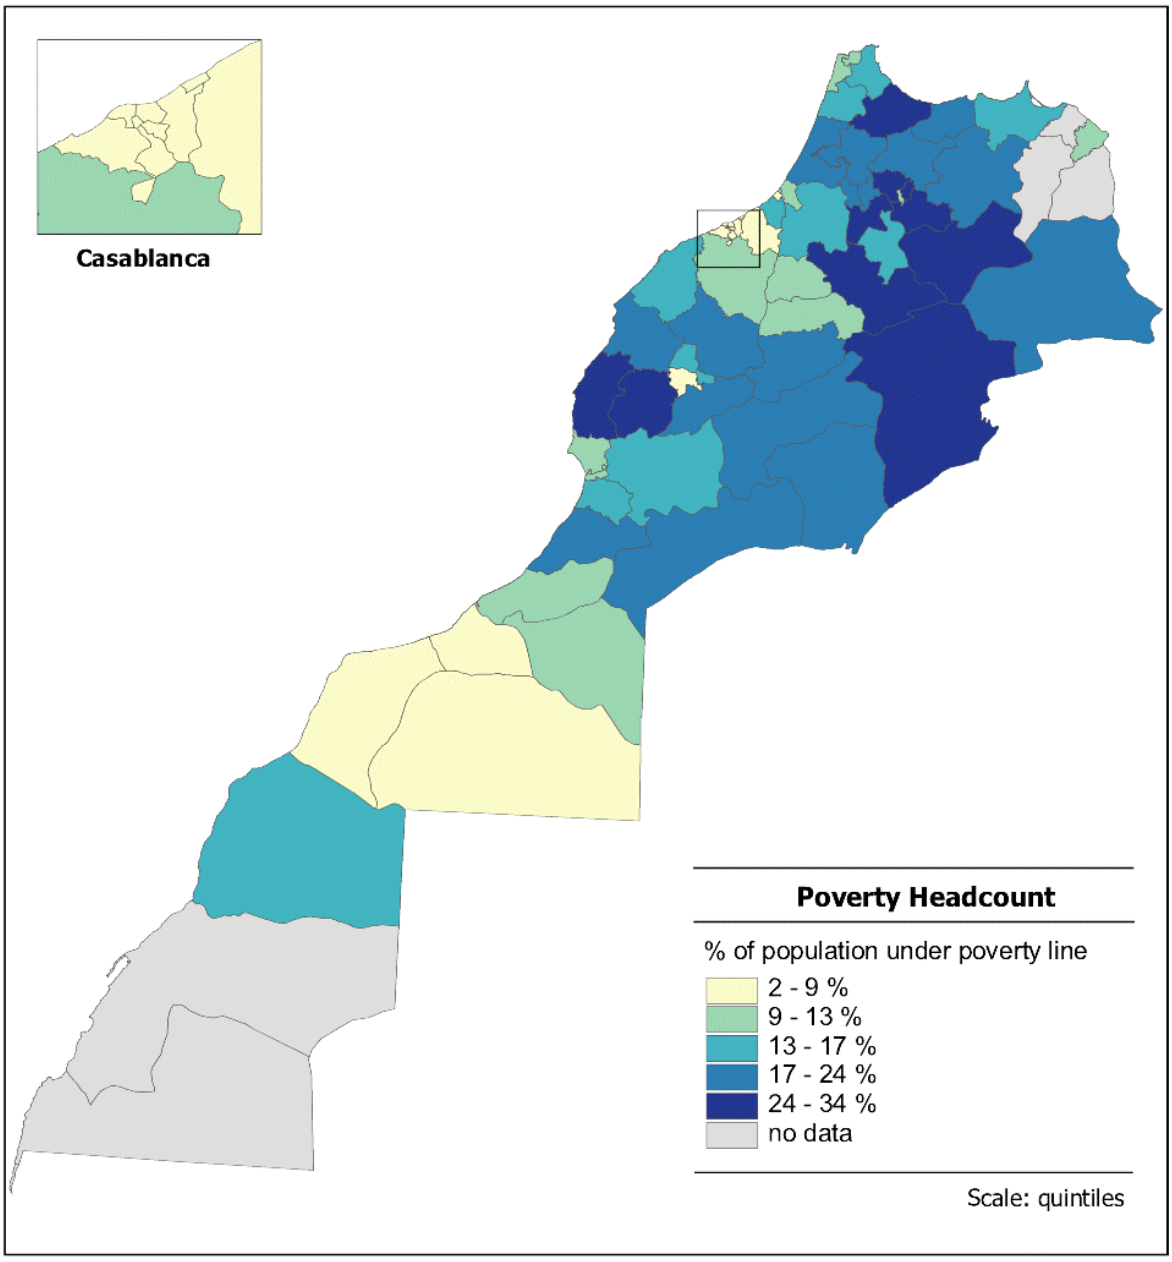

Map 2: Population Density by Province

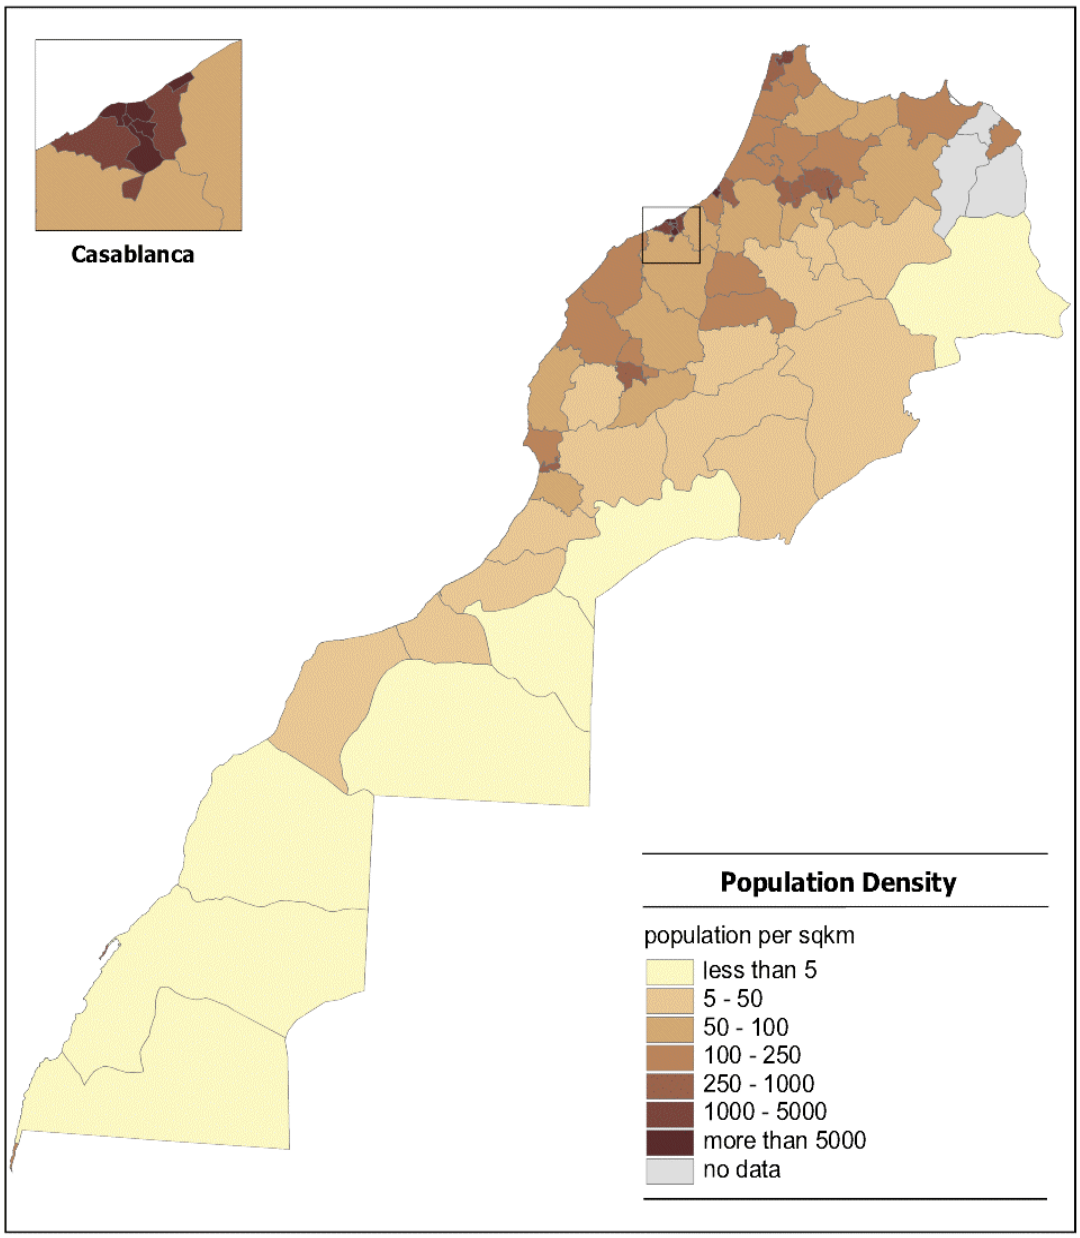

Appendix Tables 1-3 report estimates of the headcount rate (also referred to as FGT0, following Foster-Greer-Thorbecke, 1988) at the level of individual regions, provinces and communes. To guage the statistical precision of estimates at the sub-region level (provinces and communes) these Appendix tables also report standard errors on estimates. A general impression of overall precision levels can be gauged from Figures 1 and 2.

Figure 1 presents estimates of rural poverty for all communes with their accompanying 95% confidence interval. Figure 2 presents an analogous picture at the province level. It is clear that the commune and province-level poverty estimates have been produced with some imprecision. However, for a non-negligible fraction of all possible comparisons, at both the commune and provincial level, poverty has been sufficiently precisely estimated to permit meaningful comparisons.

Figure 1

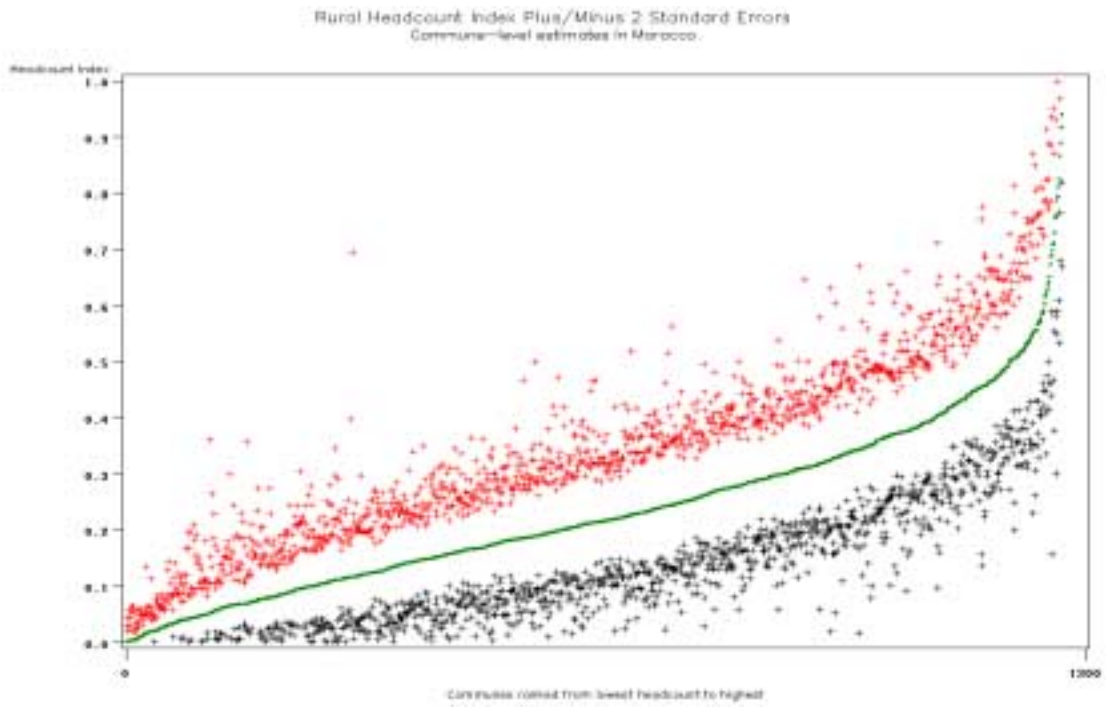

Figure 2

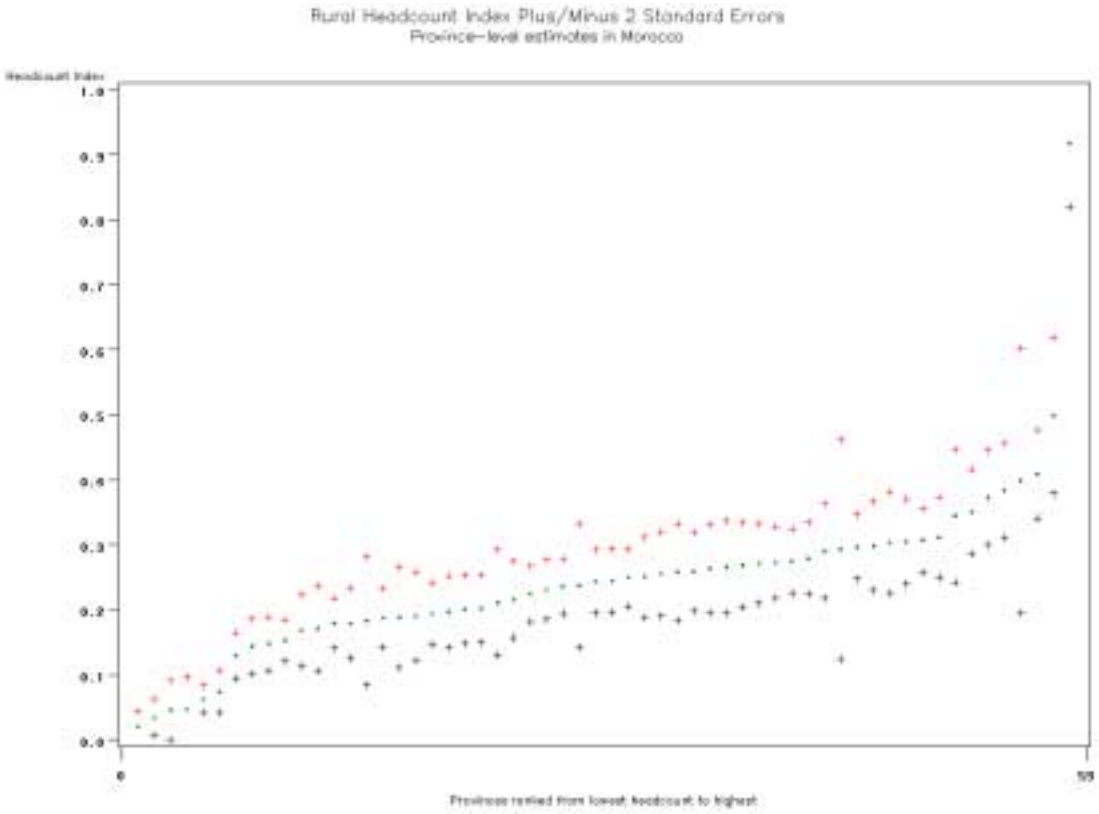

The statistical precision of poverty estimates is in some respects a function of the location of the poverty line. Figure 3 shows that when we set a poverty line 50% higher than that underpinning the estimates in Appendix Tables 1-3 and the Figures above, more commune-level heterogeneity in poverty is observed. As standard errors don't generally increase in proportion with the poverty line, a larger set of statistically significant pairwise poverty comparisons across communes are feasible with this higher poverty line.

Figure 3

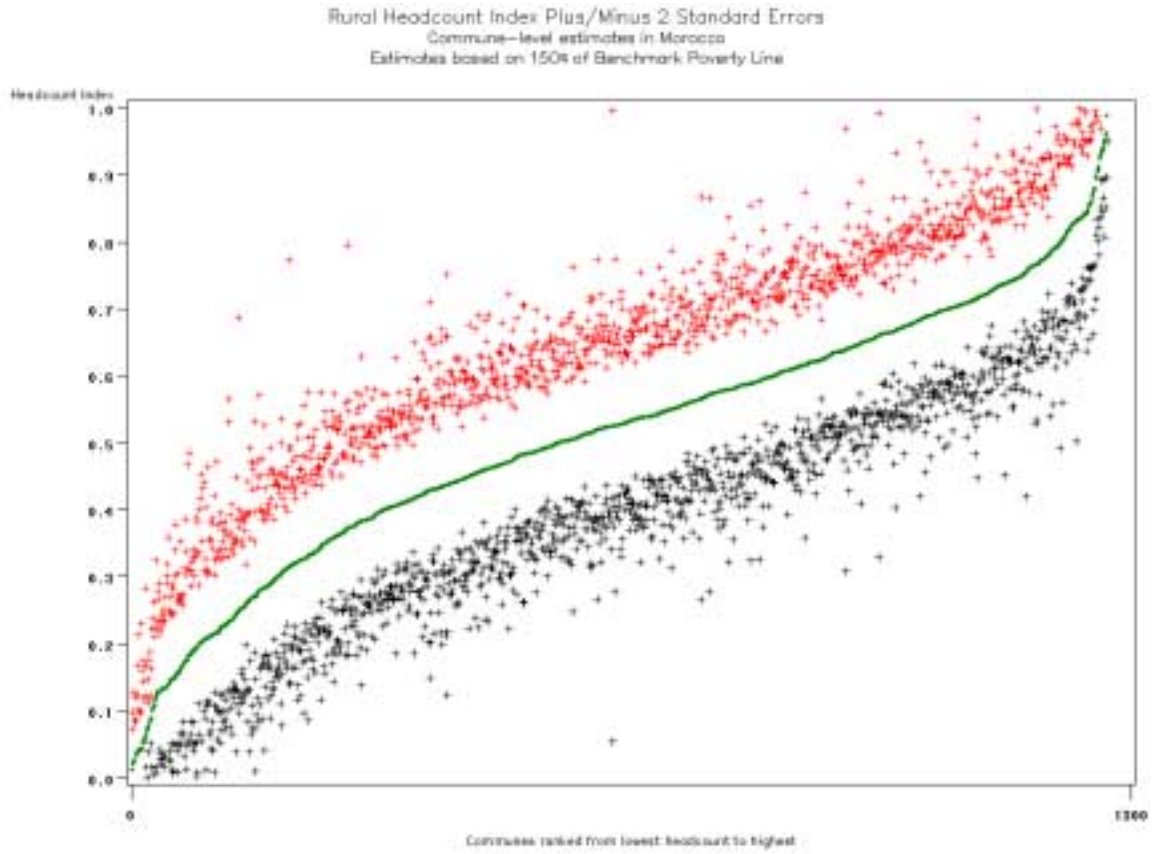

Map 1, described above, provided a glimpse of the heterogeneity in poverty that exists at the provincial level. To what extent is this heterogeneity amplified at the commune level? We have not yet succeeded in compiling the commune-level locational codes necessary to produce maps of poverty at this level of spatial disaggregation. This prevents us from directly comparing Map 1 with a similarly constructed map at the commune level. Nonetheless, Figure 4 shows that as poverty is progressively disaggregated to a more disaggregated level the distance between least poor and most poor locality widens markedly. In particular, when one moves from some 53 rural provinces to just under 1250 rural communes the heterogeneity of estimated poverty rates increases dramatically. Poverty in Morocco varies markedly across communes with some exhibiting near universal poverty (headcount rates near 100%) while in others the poverty rates are negligible.

Figure 4: Heterogeneity of Poverty by Level of Disaggregation

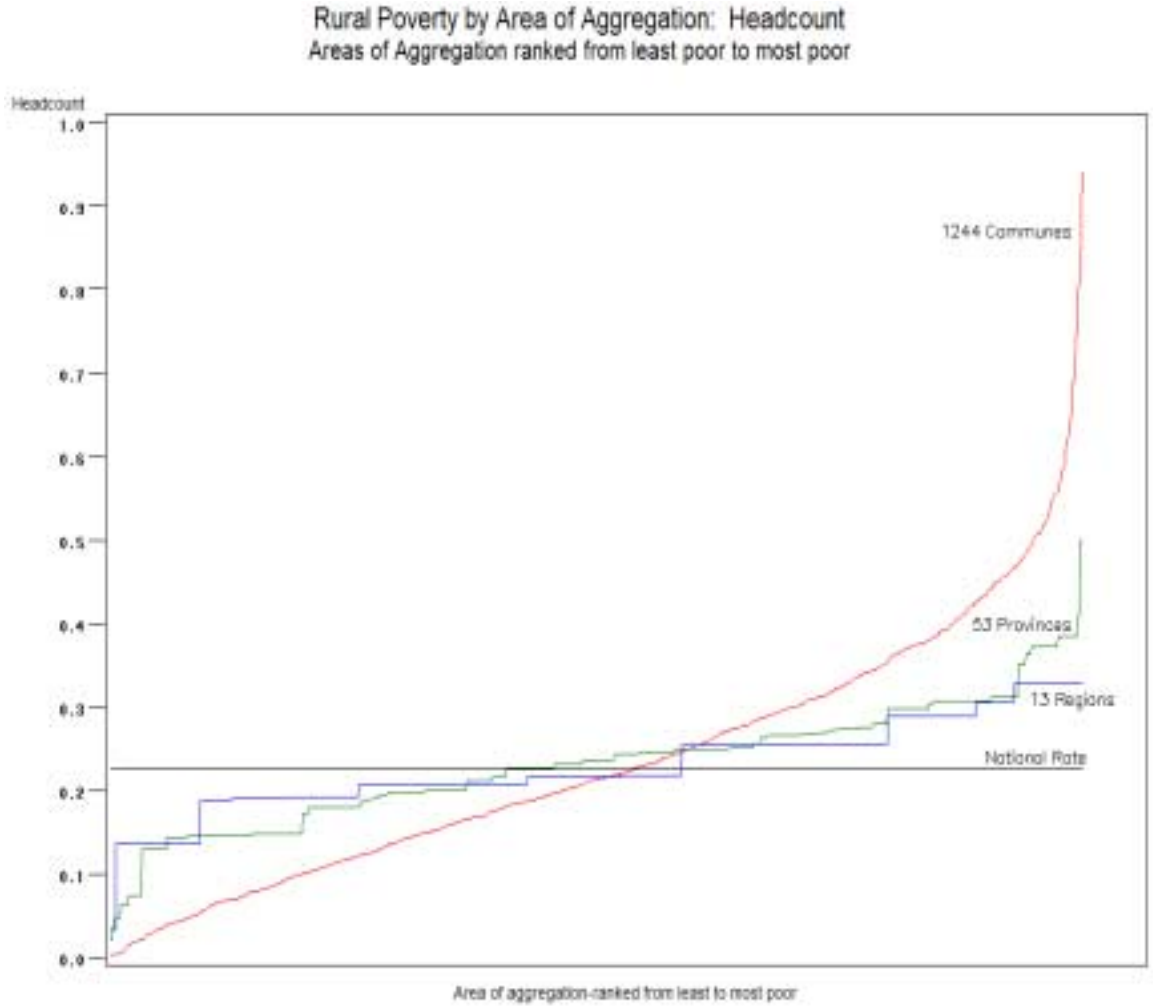

Finally, Appendix Table 4 presents commune level estimates of average per capita expenditure, inequality, and three poverty measures (the headcount rate, poverty gap and squared poverty gap) for the higher poverty lines (the benchmark lines scaled up by 50%).

The Geograpy of Inequality: Are Neighbours Equal?

The discussion above pointed to a marked heterogeneity of poverty across communes in Morocco. This heterogeneity in welfare can also be examined from the perspective of inequality. To what extent do differences between communes in average per capita consumption account for overall inequality in Morocco? How important are differences

in consumption levels *within* communities to this assessment of overall inequality? These questions can be readily investigated on the basis of a decomposition of inequality into a between-commune component and a within-commune component. The former component asks how much of overall inequality would remain if it was assumed that within communes all individuals had the same consumption level (equal to the average per capita consumption level of the commune) and hence the only variation in consumption that one observed was due to differences in average consumption levels between communes. The latter component ask the analogous question of how much overall inequality would remain if differences between communes in average per capita consumption were assumed away. Table 2 presents results from such a decomposition exercise carried out at progressively higher levels of geographic disaggregation. We can see that when inequality is decomposed by region, the dominant share of inequality is attributable to within-region inequality. As inequality is decomposed at the provincial level, the between-province share of overall inequality rises to 21% when only rural populations are considered, and 18% when rural and urban populations are combined. By the time inequality is decomposed at the commune level, the between-commune share of overall inequality is above 50% in rural areas and as high as 40% when communes are not broken down into rural and urban subgroups.

Table 2: Decomposing Inequality by Geographic Groups
(General Entropy Class Measure with parameter value of 0.5)

	Between-Group Share of Overall Inequality			No. of Geographic Units (rural and urban combined)
	Rural	Urban	Rural + Urban	
National	0%	0%	0%	1
Region	21%	3%	18%	14
Province	29%	10%	18%	67
Commune	55%	17%	40%	1517
Household	100%	100%	100%	4, 291, 792
Total Inequality (GE 0.5)	0.287	0.257	0.298	

The decomposition exercise in Morocco reveals a remarkably high degree of inequality that is attributable to differences between geographic units, particularly in rural areas. This is no doubt due, at least in part, to the tremendous variation in terrain and geography

that can be found in Morocco. The findings reported in Table 2 can be contrasted with those observed in other countries in which similar exercises have been carried out (Table 3). It is clear that Morocco stands out as a country in which geographic differences in rural areas play a particularly important role in understanding consumption inequality.

Table 3. Decomposition of Inequality Between and Within Communities

Level of Decomposition	Number of Sub-Groups	Within-group inequality (%)	Between-group inequality (%)
Ecuador			
All Communities	1579	58.8	41.2
Urban	664	76.7	23.3
Rural	915	85.9	14.1
Madagascar			
All Communities	1248	74.6	25.4
Urban	131	76.7	23.2
Rural	1117	81.9	18.1
Mozambique			
All Communities	424	78.0	22.0

Our communities in Ecuador are Zonas in urban areas and Parroquias in rural areas.
Communities are Firiasana (communes) in Madagascar and Administrative Posts in Mozambique.
Source: Elbers, Lanjouw, Mistiaen, Özler, and Simler (2004).

One of the important reasons for taking an interest in local level inequality concerns the possibility that decentralized efforts to fight poverty (via community based development efforts, for example) might be prey to capture by local elites.[3] The decomposition exercise reported in Table 2 suggests that in Morocco a sizeable portion of overall inequality is due to differences between communes in average consumption levels, and that the share of within-group inequality is generally lower than what has been found elsewhere. Does this mean that Moroccan policy makers interested in directing their efforts at poor communities can be complacent as to the possible pernicious effect of inequality within communities on their programs? Figure 5 suggests that this may not be

[3] See Elbers, et al (2004) and Mansuri and Rao (2004) for a discussion of the literature on these questions.

the case. Although much of inequality in Morocco can be attributed to differences across communes, there is some evidence that inequality within poorer communes tends to be higher than in the less poor communes. There should be no presumption that the poorest communes in Morocco are also those where differences in consumption levels across households within these communes are minor.

Figure 5

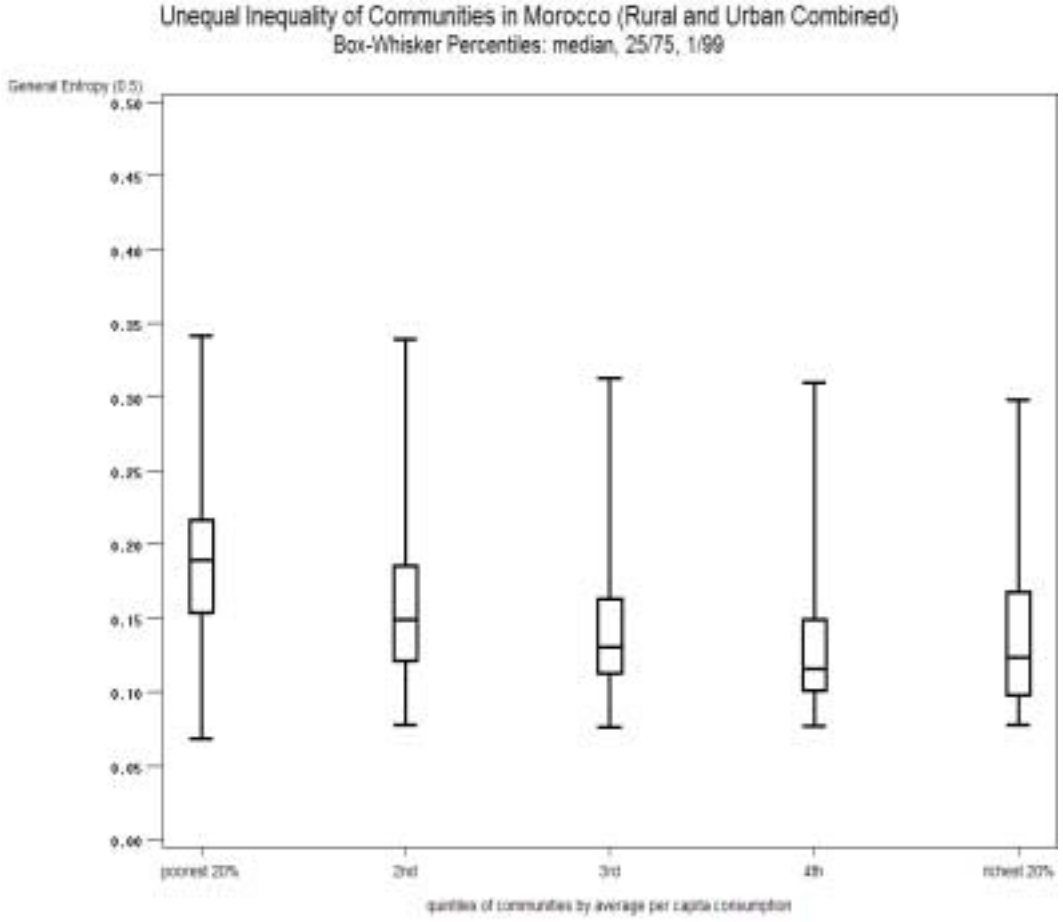

References

Elbers, C., Lanjouw, J.O. and Lanjouw, P. (2002) 'Micro-Level Estimation of Welfare' Policy Research Working Paper 2911, Development Research Group, the World Bank, Washington D.C.

Elbers, C., Lanjouw, J.O. and Lanjouw, P. (2003) 'Micro-Level Estimation of Poverty and Inequality', *Econometrica* 71(1): 355-64.

Elbers, C., Lanjouw, P., Mistiaen, J., Özler, B. and Simler, K. (2004) 'On the Unequal Inequality of Poor Communities', *mimeo*, Development Economics Research Group, World Bank, Washington D.C.

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.

Mansuri, G. and Rao, V. (2004) 'Community Based (and Driven) Development: A Review', *World Bank Research Observer*, (forthcoming)