Geographically Informed Policy Response and Intervention

Modeling the Spatial Non-Stationarity of Poverty Determinants

March 14, 2008
Addis Ababa, Ethiopia

Paul Belanger, PhD GISP
Geoinformation Systems Section
ICT, Science and Technology Division

and

Adrian Gauci
MDG / Poverty Analysis and Monitoring Section
African Center for Gender and Social Development

United Nations Economic Commission for Africa

Poverty modeling is dominated by the global, cross-country regression framework in which poverty research too often focuses on recovering previously unexplored drivers of poverty. Global regression models yield spatially invariant parameters, implying that poverty is constructed and perpetuated by ubiquitous and equally salient causal mechanisms over space. Such a conceptualization of poverty suggests that the stimuli from policy and programmatic interventions will produce the same responses everywhere. We employ geographically weighted regression to empirically illustrate that poverty is locally constructed and that parameter invariance should not be assumed. Instead, non-stationarity coupled with rich subnational data can be profitably leveraged towards geographically targeted poverty reduction interventions.

Introduction

Poverty analysts have long understood that its alleviation is predicated on understanding underlying determinants and agents. Much of the empirical literature has attempted through the classical linear regression model to assess what macroeconomic, social, geographic, and institutional factors best explain variance in levels of poverty. By extension, an understanding of these drivers may suggest a specific mixture of policy response and intervention.[1]

The literature, however, is by no means unified in its findings.[2] Some studies find that social diversity (*e.g.*, linguistic, religious, and ethnic cleavages) has some measurable impact while other studies do not. The development economics literature is likewise inconclusive on the role of trade liberalization, fiscal policy, structural adjustment, and indeed the role economic growth itself play in poverty reduction. Other analyses variously find that migration, demographic momentum, urbanization, disease prevalence, and land tenure systems are salient and covary with poverty given a set of control variables. Still others debate to what extent geography or institutions better inform our understanding of poverty.

[1] Ravaillion (1996) offers a useful summary of approaches and challenges to modeling poverty.
[2] See Collier and Gunning's (1999) review of several published poverty models.

While individual- and household-level survey data are sometimes used, it is common to many of these empirical investigations to use national-level metrics as the units of analysis. In the case of Africa, for example, some proxy for poverty is regressed on a series of exogenous variables for the 53 member States of Africa, or some subset where data are lacking. The analytical caveat to such an empirical approach—that many fail to recognize or at least acknowledge—is that both the selection of the state as the unit of analysis and the specification of a global regression model may mask significant intra-national variance. It may, moreover, average away and render locally important relationships insignificant in a global model of the continent.

As such, the present paper argues for a critical reexamination of the conventional assumption of spatial stationarity in the agents of poverty. Is the relationship between poverty and some environmental, social, institutional, or geographic condition spatially invariant? If we find from a global model of African poverty that linguistic diversity, for example, has a non-zero, positive, and significant effect, can and should we assume that that relationship holds, with statistical effect, across the continent or at the subnational level? Conversely, can we assume that because a potential poverty determinant is found to be an insignificant predictor in a cross-country model that it is equally irrelevant across the continent?

Our intent here is therefore not to reengage the debate over model specification *per se* but, rather, to demonstrate that a sensitivity to local forms of analysis can deepen our understanding of poverty. Our purpose is not to introduce and test for the salience of particular poverty determinants—though much work remains to be done here—but instead to explore for spatial non-stationarity in poverty determinants that are well known and theoretically informed.

Following a review of some of the empirical literature on African poverty, geographically weighted regression is then introduced as an innovative alternative to global model specification that allows us to unpack and tease out variations in the local construction of poverty. We next demonstrate how a local model of poverty determinants yields a more nuanced understanding of its construction and, as such, suggests locally tailored poverty reduction strategies and interventions. Concluding that we must not assume spatial invariance but, rather, explicitly test for non-stationarity, a potentially fruitful research trajectory is presented. We commence, though, with a review of the empirical literature on African development.

Geographic 'Determinants' of Poverty

The empirical literature is replete with analyses that attempt to tease out *the* determinant or set of explanatory variables that account, *ceteris paribus*, for the continent's disproportionately high levels of poverty (Collier and Gunning, 1999). That is, in virtually all cross-country poverty regressions, there remains an unexplained factor or set of factors that can only be captured through an African dummy variable.[3] It is found to be large and significant (Barro and Lee, 1994; Easterly and Levine, 1997; Acemoglu *et al.*, 2001b). Collier and Gunning (1999: 65) note that "…slow growth [in Africa] is explicable in terms of a distinctive effect of variables in Africa, which shifts the question to explaining this different response."

To eliminate that seemingly inexplicable African factor, the poverty and growth literature has expanded considerably and has taken aim at explicating and testing for the theoretical traction and empirical efficacy of a host of hypothesized poverty determinants. Recurrent themes, amongst many, include demographic characteristics and

[3] Collier and Gunning (1999) provide an extensive and accessible summary of the relevant literature.

momentum (Bongaarts, 1994), the natural environment and resource endowments (Diamond, 1997; Hibbs *et al.*, 2004), the role of institutions,[4] social capital, socioeconomic cleavages (Easterly and Levine, 1997; Collier, 1998), and the effects of trade and financial liberalization (Sachs and Warner, 1997; Serieux and McKinley, 2008).

The focus in much of the literature cited above has been towards identifying either missing or mis-specified variables in propelling growth and reducing poverty. The academic growth literature has embraced the growth-regression framework stemming in part, as Francisco Rodriquez (2007: 1) observes, because of "the inherent appeal of finding 'causes' of growth that can serve as magic bullets in the development process." Van de Walle (2001: 14), for example, asserts that "political institutions hold the explanatory key to the African crisis and there will be no successful economic reform without prior reform of the region's politics."

Even if one could accept such a sweeping generalization, can one accept that underdeveloped institutions—or for that matter any poverty determinant—are a causal agent of poverty with equal effect across the continent? Some social scientists have and are thinking beyond the missing variable strategy to modeling poverty and have begun, instead, to question the resulting one-size-fits-all poverty model achieved from such studies. Conventional cross-country regression approaches are increasingly seen to limit our understanding of poverty. Hentschel *et al.* (1998: 2), for example, observe that:

> [t]he empirical relationship between poverty or inequality and indicators of development, such as economic growth, is typically examined in a cross-country regression framework. It is difficult, however, to control for the enormous heterogeneity

[4] For an engaging debate on the role of geography *v.* institutions in the construction and perpetuation of poverty, interested readers should consult Acemoglu *et al.* (2001a), Acemoglu *et al.* (2001b), Sachs (2001), Bloom and Sachs (1998), Sachs (2003a and 2003b), Woods (2004), and Hibbs *et al.* (2004).

> which exists across countries; heterogeneity which may mask true relationships.

Likewise, and more concretely, Coulombe and McKay (1996: 1016) note in their study of poverty determinants in Mauritania that:

> Mauritania is different in many ways…from other African countries in which poverty has been studied. These differences, including its economic structure and its geographic and socioeconomic characteristics, may mean that poverty is likely to be different in nature from that in other African countries.

Rodriquez (2007: 2), while taking *linear* growth regressions to task, identifies a more general shortcoming of the cross-country regression approach:

> [t]he foremost problem is dealing with real world complexity. The workhorse growth regression embodies a particular vision of the world that assumes, implicitly, that the same model of growth is true for all countries.

We believe, further, that this framework largely obscures local environmental, social, political, and economic processes underpinning poverty and renders them difficult to empirically recover. So, too, does Fofack (2000: 214) appreciate that the construction of poverty is scale dependent and place specific:

> [t]he causes and determinants of poverty…are variable. At the aggregate level differences in the potential for income-generating activities and wage inequality may constitute important factors; at the regional and district levels human capital, access indicators, and location of infrastructure may be more critical.

Analytically, Rodriquez recognizes (2007: 2) that changes in an explanatory variable are:

> assumed to have the same effect in a poor country as in a rich country, in a primary-resource exporter as in a manufactures exporter, and in a country with well-developed institutions as in a country with underdeveloped institutions.

The challenge that Fofack, Rodriquez, and others have identified is that the construction of poverty is spatially variable and scale dependent. Cross-sectional regression models yield a unitary set of coefficients that, after controlling for other factors, measure the impact of an exogenous variable on the level of poverty. They are, however, global average parameters and as such may mask local variations in the rate of stimulus-response. In many ways, this challenge represents the spatial analogue to Simpson's classic paradox.

Simpson's paradox is a statistical paradox in which the relationship between X and Y becomes apparent, is eliminated, or even reversed upon the introduction of a confounding covariate Z (*i.e.*, associations in aggregate form may disagree when disaggregated). While Simpson's original note and most subsequent demonstrations of the paradox employ a non-spatial Z covariate, spatial covariates may also reveal the paradox (Simpson, 1951; Knapp, 1985; Wagner, 1982; Appleton *et al.*, 1996).

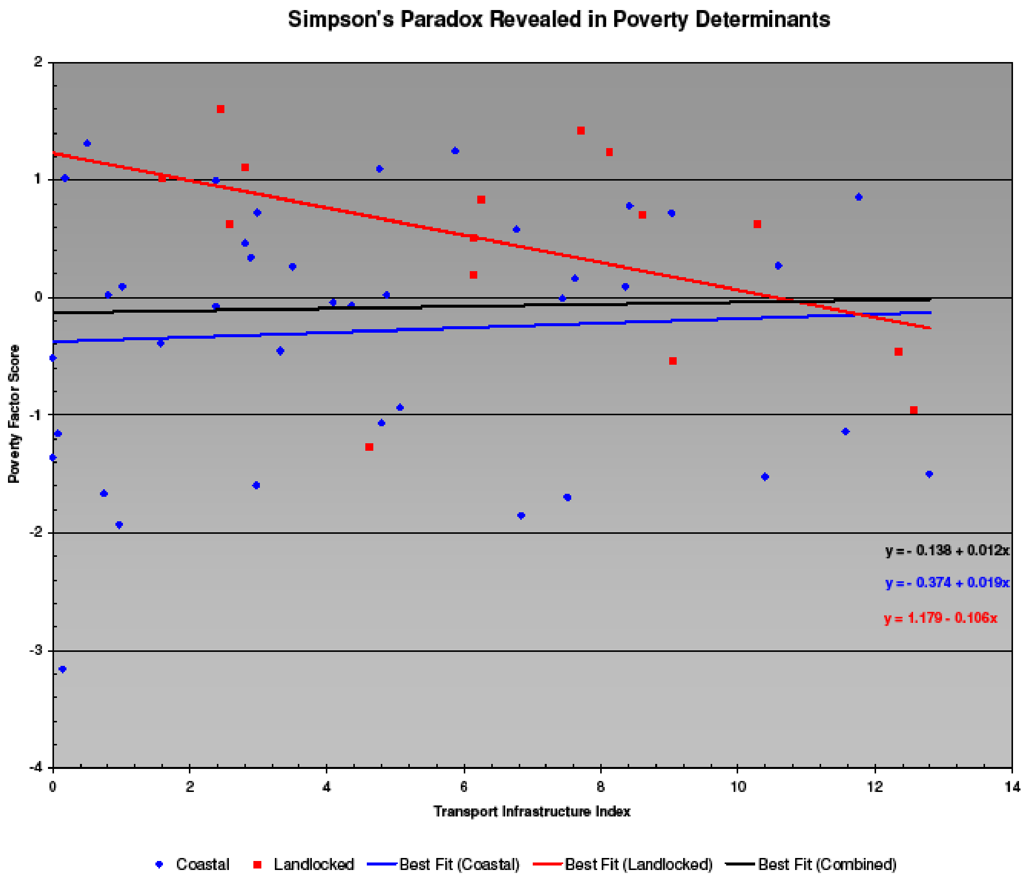

Figure 1: Example of Simpson's Paradox in Poverty Determinants

Figure 1, for example, illustrates that a potential relationship between the level of transport infrastructure[5] and poverty[6] may exist when we disaggregate the countries of Africa by their coastal or landlocked status. In aggregate form, as shown by the combined line of best fit, there is no evident relationship between transport infrastructure and poverty with a coefficient of 0.012 being close to zero. The disaggregate data and trend line, however, reveal a potential relationship in landlocked countries between poverty and density of the transport network and no apparent relationship in coastal states. In a typical cross-country poverty model we would, on the basis of these aggregate data, conclude that transport infrastructure is not a salient predictor of poverty. This spatial variant of Simpson's paradox highlights the risk of planning poverty reduction programmes and interventions on the basis of aggregate data.

This simple example demonstrates that a global coefficient may not, in fact, be stable over the study area and here we have captured spatial non-stationarity (for the purposes of illustration) through a rather crude binary spatial variable of access to the sea.[7] If we conceptualize the relationship between poverty and transport infrastructure as a parameter surface, then global cross-country regressions yield a planar surface. By then introducing the landlocked status of a country we can envisage a bi-level surface of the poverty-transport relationship, with coastal states on one plane and landlocked countries on another.

If, however, we extend this approach to its logical end, it yields the possibility that local levels of poverty as a function of transport infrastructure can be represented not by a unitary plane, or by a

[5] Proxied here by road network density as kilometers of road per square kilometer.
[6] Measured here as a principal component comprised of GDP per capita, infant mortality, and child malnutrition.
[7] We could capture this effect by introducing an interaction term between landlocked status and the transport infrastructure index.

bifurcated plane reflecting landlocked status but, rather, by a continuous parameter surface upon which the stimulus-response rate is non-stationary. In the next section we present a recently developed technique to detect non-stationarity and generate spatially variable parameter surfaces.

Spatial Non-Stationarity and Geographically Weighted Regression[8]

In the standard linear regression model:

$$y_i = \beta_0 + \beta_1 x_{1i} + \beta_2 x_{2i} + \ldots + \beta_n x_{ni} + \varepsilon_i \qquad \textbf{(1)}$$

y_i is a measure of poverty at location i, calculated as an additive function of a global intercept β_0, a local stochastic error term ε_i, and the product of n global parameters and locally measured exogenous variables $\beta_1 x_{1i} + \beta_2 x_{2i} + \ldots + \beta_n x_{ni}$.

Equation 1 implies not only that the parameters remain constant over space but also that all observations contribute equally to their calibration at location i. Geographically weighted regression, a variant of weighted regression that accounts for the spatial distribution of observations, allows for the estimation of location-specific parameters:

$$y_i = \beta_0(e_i, n_i) + \beta_1(e_i, n_i)x_{1i} + \beta_2(e_i, n_i)x_{2i} + \ldots + \beta_n(e_i, n_i)x_{ni} + \varepsilon_i \qquad \textbf{(2)}$$

Letting (e_i, n_i) represent the easting and northing of location i, poverty at i is calculated as an additive linear function of a place-specific intercept $\beta_0(e_i, n_i)$ and the sum of locally measured exogenous variables multiplied by their place-specific coefficients $\beta_1(e_i, n_i)x_{1i} + \beta_2(e_i, n_i)x_{2i} + \ldots + \beta_n(e_i, n_i)x_{ni}$.

[8] Space limitations here enable us to present only the essentials of geographically weighted regression (GWR). A full and highly accessible treatment of GWR can be found in the seminal work of Fotheringham *et al*. (2002).

Deriving a continuous parameter surface from punctiform observations in geographically weighted regression requires that we define an optimal bandwidth around each regression point i to limit observations included in the model's calibration at i. We could, for example, impose an arbitrary or theoretically informed kernel of some number of observations or of some distance.[9] This, in and of itself, will produce continuous parameter surfaces (akin to a moving window regression over space). However, the innovation in geographically weighted regression is, as its name implies, to weight the observations around i given their distance from i.

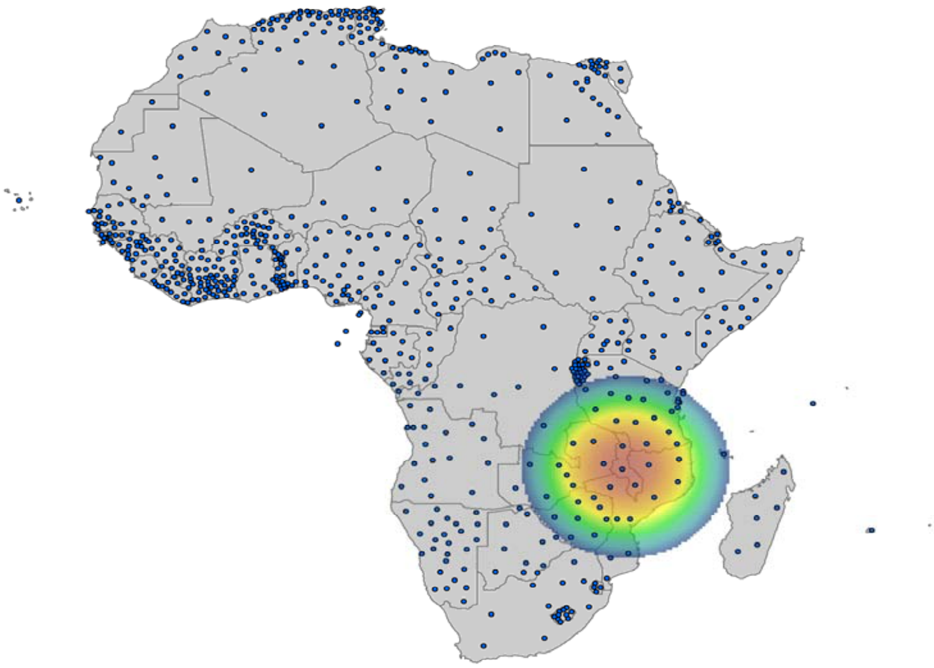

Figure 2: Spatial Kernel (in 2D)

A spatial kernel is imposed around each location i such that observations closer to i have more weight than observations distant from i. Doing so operationalizes Tobler's (1970: 236) first law of geography that *everything is related to everything else, but near things*

[9] Beyond the scope of this paper, the choice of bandwidth considerably impacts the calibration of a geographically weighted regression model. On bandwidth selection and the choice of fixed or adaptive kernels, readers are encouraged to consult *pp*. 44-51 and *pp*. 59-62 in Fotheringham *et al.* (2002).

are more related. These spatial weights thus enable a more locally sensitive estimation of the model in the vicinity of location i instead of imposing a global average set of coefficients.

Figure 2 depicts a spatial kernel for, say, estimating a poverty model centred on Malawi. Distant observations beyond a particular kernel size, in this case a 1200km bandwidth, have a weight of zero and do not impact the estimation of local poverty factors. Observations from Guinea, Algeria, or Djibouti, for example, would have no influence on the estimation of local poverty factors in Malawi as illustrated in Figure 2.

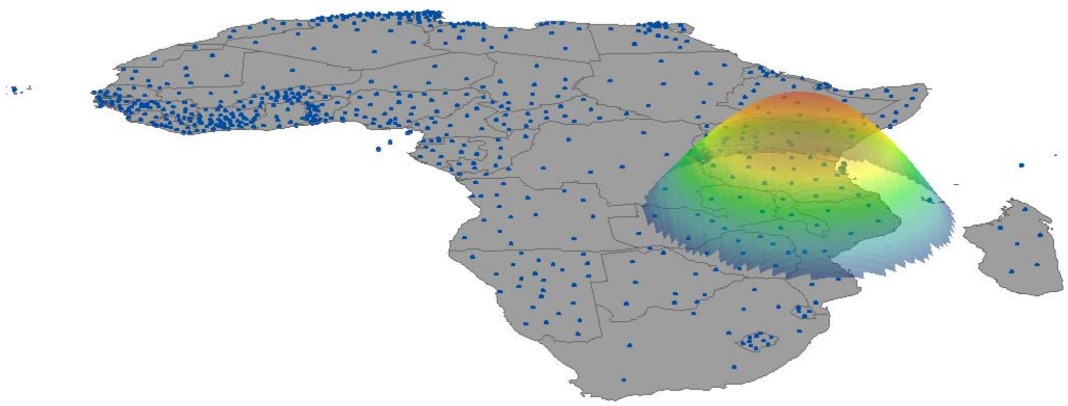

Figure 3: Gaussian Spatial Kernel (in pseudo 3D)

Even within the kernel, the estimation of relationships between poverty and its determinants in Malawi, for example, are more influenced by observations of poverty and its determinants in Malawi, Zambia, and Mozambique than they are by observations from Zimbabwe, Botswana, or the DRC as depicted in Figure 3. Though still captured by the kernel these latter data points lie at the kernel's fringe and are more heavily discounted given their distance from the regression point. The weighting scheme is fully flexible and as in standard weighted regression it is possible to specify a custom spatial weighting function.[10] As portrayed in Figure 3 a standard Gaussian distance decay function was used in this case:

$$W_{ij} = \exp(-0.5(d_{ij}/b)^2) \qquad \textbf{(3)}$$

where W_{ij} represents the weight between regression point i and observation j, d_{ij} is the distance between i and j, and b—the bandwidth—is the size of the kernel.

A brief review of the relevant literature on poverty modeling suggests some recognition that the global cross-country regression framework fails to capture locally salient poverty determinants. As in social science more generally—and unlike the physical sciences—there may not exist a universally generic poverty model. Rather, local differences in socioeconomic composition, demographic characteristics, economic structure, and geography may contextually combine in different ways in the formation and perpetuation of poverty. The structural relationships between poverty and its determinants may be intrinsically different over space such that fixed interventions and stimuli produce variable and perhaps counterproductive responses. Can this claim be empirically sustained? Can we assume spatial stationarity in the construction of poverty? Equipped with geographically weighted regression we seek to explore these questions in the next section.

Analysis and Results

While the relevant literature would appear to concur that poverty is a multi-dimensional phenomenon,[11] most empirical studies revert to

---

[10] For a detailed discussion of weighting functions, interested readers should consult *pp*. 56-59 of Fotheringham *et al*. (2002).

uni-dimensional income- or expenditure-based metrics to quantify poverty (Maasoumi and Lugo, 2008). In this paper, we utilize three (3) indicators of poverty: per capita income, infant mortality, and child malnutrition. Income, as already mentioned, has been used in much empirical work, is widely available, and is theoretically tractable. Infant mortality—notwithstanding some collinearity with income—taps unique health-related dimensions of poverty such as access to health care facilities, medicine, and physicians (Coulombe and McKay, 1996). Likewise, child malnutrition captures elements of dietary nutritional content and caloric intake that are not directly related to income (Sahn and Stifel, 2002). By reducing these three variables through factor analysis to a single component—a composite measure of poverty—we derive a more robust metric of poverty and hopefully mitigate an overreliance on a single variable.

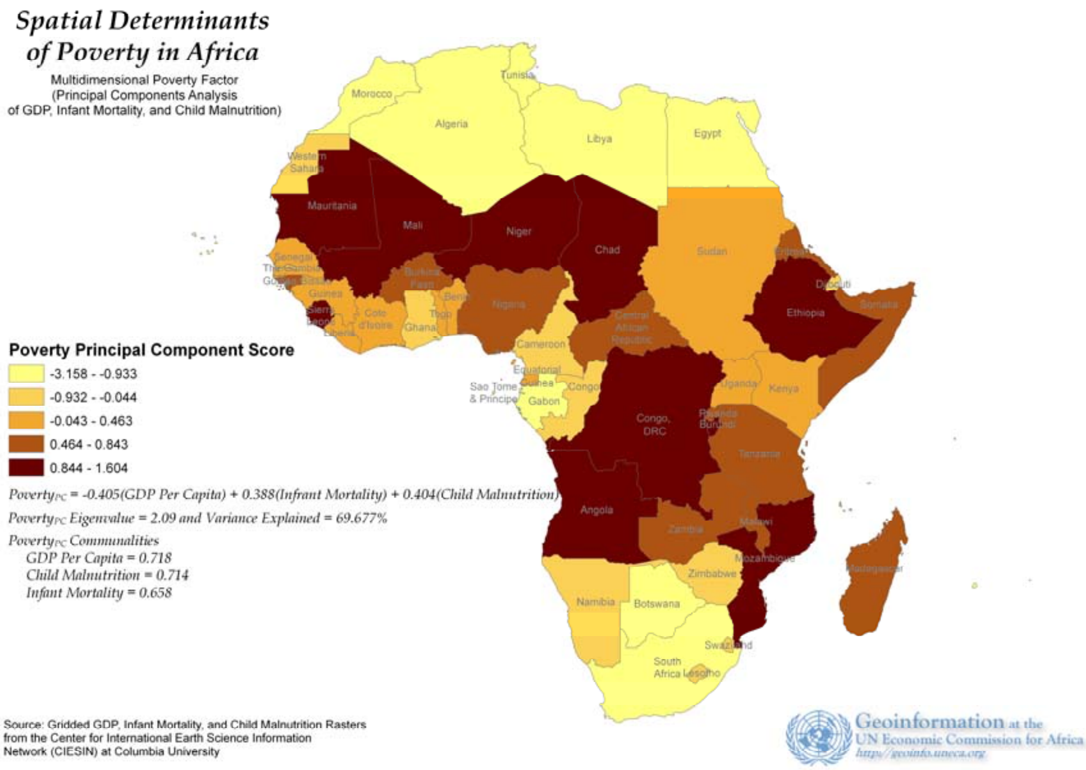

Figure 4: Africa Poverty Princicpal Component

---

[11] From Sen's (1985, 1987) oft-cited admonition of income-based poverty to more recent (McKinley, 2006) critiques, income-centric metrics of poverty are increasingly suspect.

Summarized in Figure 4 is the factor analysis reducing income, infant mortality, and child malnutrition to a composite poverty component. The factor score accounts for some 70% of the original variation in the three indicators and the loadings are relatively equal for each of the three variables. As expected, the poverty score is positively related to infant mortality and child malnutrition and inversely related to GDP per capita. Geographically, the resultant mapping of the poverty factor would appear to resonate with the conventional understanding of poverty across the country: north Africa, South Africa, and Botswana fare relatively well while the Sahelian countries, central Africa, Ethiopia, and post-conflict Sierra Leone constitute the bottom quintile. While the poverty score itself is difficult to interpret we believe that this composite score somewhat insulates our subsequent analysis from the idiosyncrasies and vagaries of a uni-dimensional poverty metric and, as such, better captures the spatial variation of a multi-criteria conceptualization of poverty.

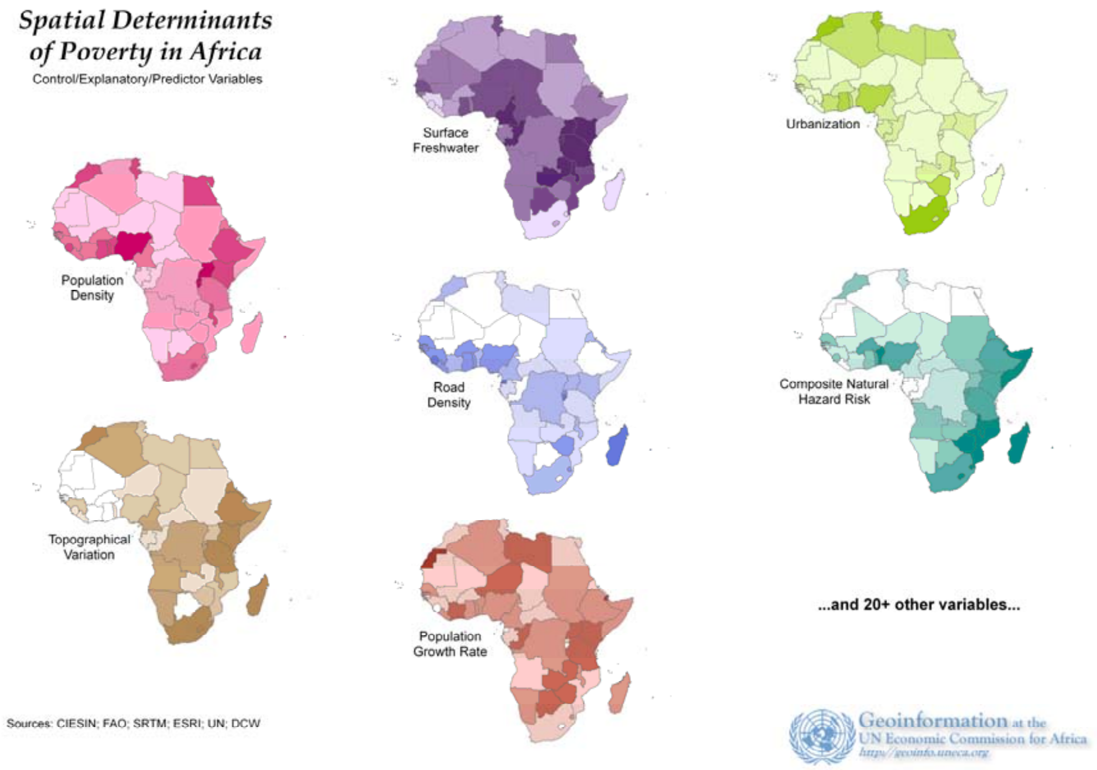

Figure 5: Potential Poverty Determinants

As our goal here is to demonstrate the need for a sensitivity to the local, our selection of exogenous variables is admittedly expedient, predicated on data availability and on theoretical rationales already developed in the relevant literature. However, additional to these conventionally studied poverty determinants, we introduce composite natural hazard risk as a predictor of poverty (including floods, cyclones, landslides, earthquakes, and drought). Owing to space limitations we map only selected predictor variables (each colour class is approximately one quintile) in Figure 5.

Table 1: Global Regression Parameter Estimates

Predictor	Coefficient	Std. Error	t*
Intercept	0.228	0.755	0.302
Demography/Settlement			
Population Density (person/km^2)	-0.000	0.002	-0.049
Pop. Growth Rate, 1960-2000	-0.002	0.002	-0.958
% Urban	-1.802	0.756	**-2.384**
Environmental			
Mean Elevation (in metres)	-0.000	0.001	-0.525
Topography (σ in metres)	0.000	0.001	0.302
Composite Natural Hazard Index	0.030	0.014	**2.020**
Agricultural Resources			
% Surface Freshwater	0.013	0.007	**1.828**
Mean Rainfall Runoff (mm/annum)	0.001	0.000	1.382
% Regosols	0.040	0.021	**1.956**
% Yermosols	-0.024	0.014	**-1.643**
Infrastructure			
Road Network Density (km/km^2)	-0.072	0.036	**-1.979**
Land Use			
% Shrubland/Savannah	0.020	0.013	1.520
% Cropland	0.012	0.013	0.909
% Bare Soil	0.006	0.008	0.683

* $t_{\alpha=0.05,n=30+} = 1.697$ are in bold; R^2 = 0.412; N = 54

A cross-country global regression would yield results shown in Table 1. A standard interpretation would include recognition that urbanization would seem to be inversely related to poverty levels, natural hazard risk is directly related, regosols also vary positively[12]

[12] Regosols are often found alongside other young or poorly developed soils in arid, degrading, or eroding areas. Regosols in desert areas have minimal agricultural significance. The low moisture holding capacity of these soils calls for frequent applications of irrigation water; sprinkler or trickle irrigation solves the problem but is rarely economic. As such, regosols are sometimes used in capital-intensive irrigated farming but the most common land use is low volume grazing.

while yermosols[13] are inversely related, and higher transport infrastructure densities are associated with lower levels of poverty. Population density and growth, elevation and surface topography, runoff, and land use are all found to be insignificant predictors of poverty in this global model. The model nonetheless performs relatively well with an R^2 of 0.41.

Given that global regressions produce invariant coefficients, a map of the model's most significant predictor--% urbanization—reveals no spatial variability in its influence, as illustrated in Figure 6. With a constant standard error the parameter is assumed to be an equally salient determinant of poverty across the continent, again as illustrated in Figure 6.

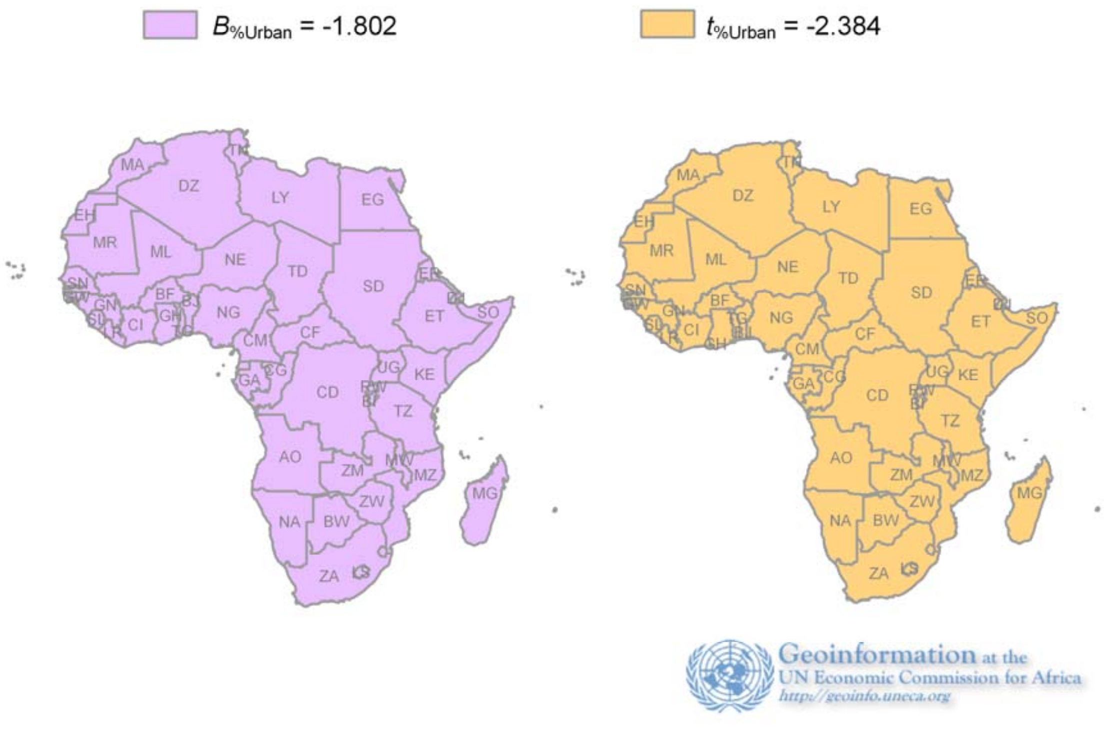

Figure 6: Spatial Variation in % Urbanization Coefficient and Significance

Interpretations such as these are required in a global regression framework as only average coefficients, standard errors, and

[13] Present in much of Africa, except for central Africa, yermosols have a wide variety of agricultural uses, though climate, topography, shallowness, or stoniness, may pose restrictions on land use. They are used for (mixed) arable farming and also as grazing land.

significance statistics are estimated. Variables found to be globally insignificant are assumed to be insignificant everywhere; variables found to be salient are assumed to be salient everywhere; those found to vary directly with the dependent variable are assumed to vary directly with it everywhere and vice versa; and the magnitude of the effect is also necessarily presumed to remain constant throughout the study area. And not only is the parameter surface planar for every determinant but the model is also assumed to explain poverty equally well (or poorly) across the continent as shown in Figure 7.

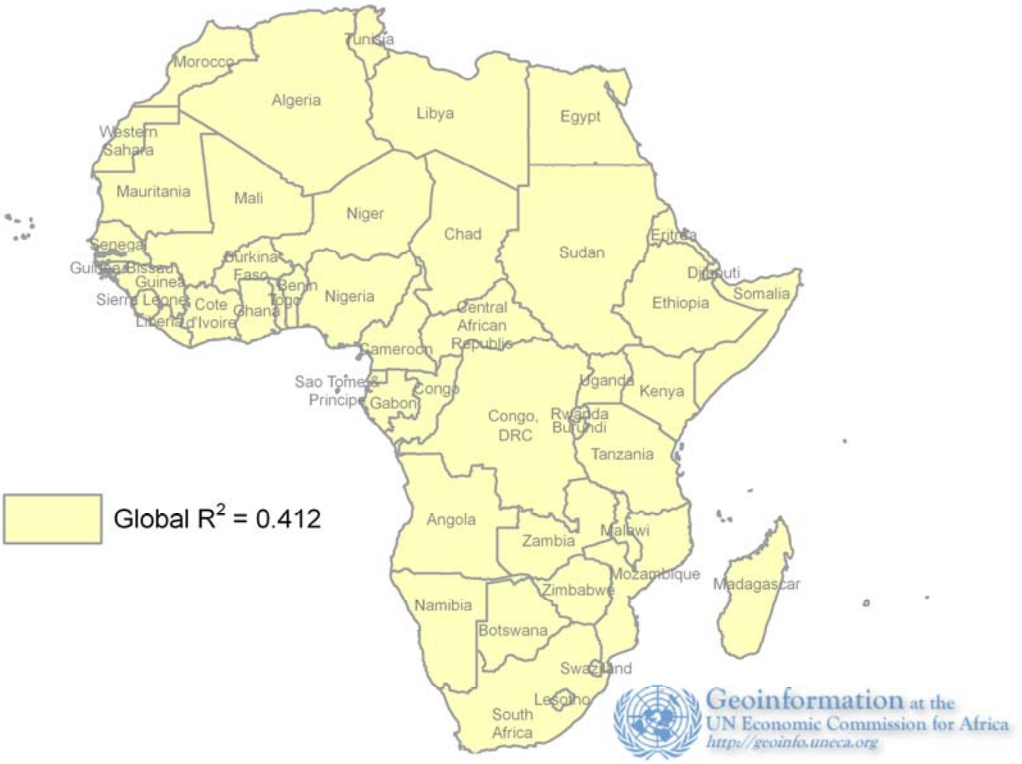

Figure 7: Global Adjusted R^2

Recasting our model within a geographically weighted regression framework reveals some interesting findings. While the objective function for bandwidth selection includes almost all observations in each local regression (*i.e.*, the closest 52 of 54 observations), the Gaussian weighting function teases out some compelling evidence of spatial non-stationarity in the construction of poverty. Table 2 summarizes the distribution of parameter values for each determinant, giving the minimum, 25th percentile, median, 75th percentile, and maximum values. The last column gives the p-value for a Monte Carlo test of the spatial variability of the coefficients for each predictor.[14]

Table 2: Geographically Weighted Regression Parameters

Predictor	Min.	Lower Quartile	Median	Upper Quartile	Max.	Spatial Variability (p-value)
Intercept	-0.977	-0.482	-0.313	-0.126	0.733	0.77
Demography/Settlement						
Population Density (person/km^2)	-0.005	-0.004	0.002	0.008	0.008	0.01
Pop. Growth Rate, 1960-2000	-0.004	-0.003	-0.002	-0.002	-0.002	0.87
% Urban	-3.752	-3.567	-1.712	-0.712	-0.409	0.00
Environmental						
Mean Elevation (in metres)	0.000	0.000	0.000	0.000	0.001	0.63
Topography (σ in metres)	-0.002	-0.002	-0.002	0.000	0.001	0.06
Composite Natural Hazard Index	0.018	0.028	0.038	0.051	0.060	0.19
Agricultural Resources						
% Surface Freshwater	0.005	0.007	0.008	0.010	0.011	0.82
Mean Rainfall Runoff (mm/year)	0.000	0.000	0.000	0.001	0.001	0.42
% Regosols	0.013	0.025	0.054	0.075	0.077	0.08
% Yermosols	-0.039	-0.037	-0.035	-0.013	-0.003	0.35
Infrastructure						
Road Network Density (km/km^2)	-0.109	-0.038	-0.022	0.000	0.083	0.57
Land Use						
% Shrubland/Savannah	-0.001	0.001	0.008	0.027	0.032	0.00
% Cropland	-0.031	-0.028	-0.016	0.023	0.027	0.00
% Bare Soil	0.000	0.004	0.012	0.017	0.023	0.20

Local R^2 Range: 0.742 – 0.868; N Nearest Neighbours = 52

In lieu of interpreting the entire range of parameter estimates for every predictor, we here observe some more insightful findings achieved through GWR otherwise masked in a global regression. Composite natural hazard risk, as suggested by the global model, varies positively with poverty across the continent; however, the effect is not uniform having more than three times the impact towards poverty construction in the most influential locales than in the least salient parts of the continent (*cf.* 0.018 *v.* 0.060). Similarly, the proportion of surface area that is fresh water uniformly varies with poverty with a ratio greater than 2 between the highest and lowest parameter estimates. Likewise, the ratio of parameter variation in regosols is

[14] Separate tables could be given showing the standard errors and t-stats for each predictor but space limitations prevent us from doing so. Instead, we map these statistics for selected predictors.

close to 6 and approximately 13 between the maximum and minimum yermosol coefficients.

Relying on global parameter estimates not only conceals this considerable non-stationarity in parameter estimates but also may not represent the relationship with poverty in any part of the study area. The coefficient for the proportion of land under crops, for example, is estimated at 0.012 by the global model whereas the GWR analysis reveals that the effect ranges from -0.031 to +0.027 and that this spatial variability is highly significant.[15] One needs to question then how well the global statistic captures the relationship in any part of Africa.

As surfaces, the parameters are easily mapped and the spatial variation in the coefficients readily gleaned. Figure 8 maps the parameter surface for the percent urban population with the geographically weighted regression surface intersected, for illustration, by the global average plane from Figure 6.

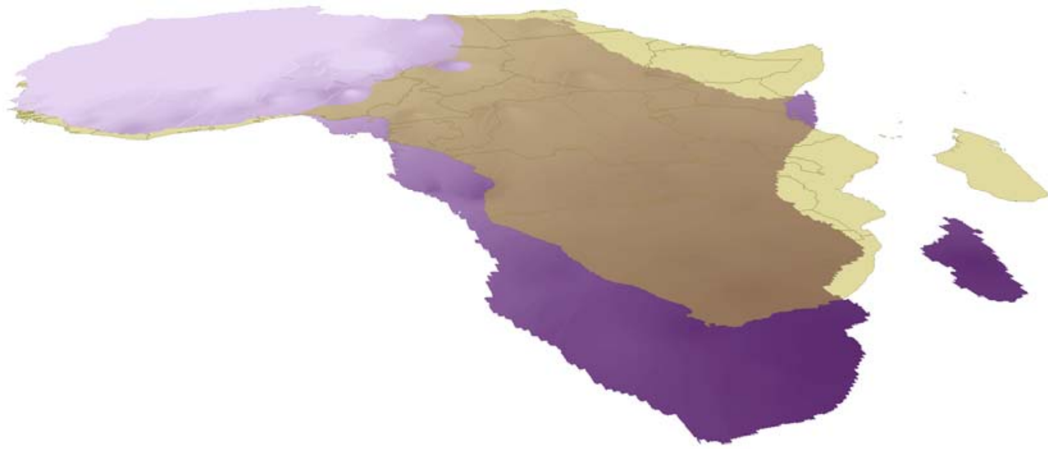

Figure 8: Percentage Urban GWR and Global Parameter Surfaces

[15] Note, though, that while the spatial variation in the parameter estimates may be statistically significant under the Monte Carlo test, one needs to examine the local standard errors to determine if the predictor differs significantly from zero at that locale.

The effect of urbanization on poverty levels is considerably more pronounced in southern, eastern, and central Africa. In much of north and west Africa, the coefficient is closer to and less negative than the global average parameter. The urbanization effect on poverty levels would appear to be considerably more influential in southern and eastern Africa, a nuance concealed by the global estimate. The surface in Figure 9 depicts the significance of the local coefficient estimate, with statistical significance achieved in much of east and southern Africa, an attenuation of it along an axis extending northeasterly from Gabon to Egypt/Sudan, and an insignificant relationship between urbanization and poverty across much of west and north Africa.

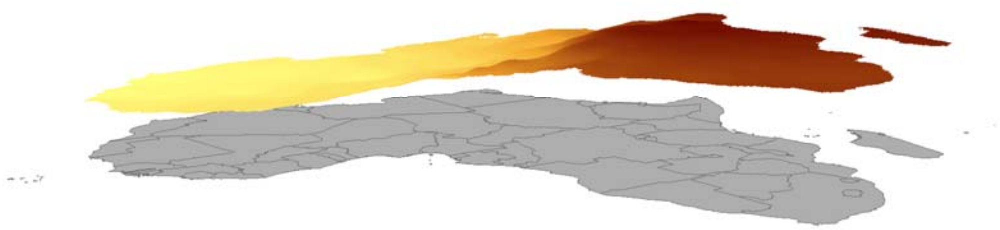

Figure 9: Urbanization Significance Surface (t-stat)

One could (and should) successively interrogate the parameter surface for each poverty determinant for non-stationarity. Is the sign (*i.e.*, the direction of the relationship) stable through the full range of estimates? Does the distribution of coefficients, particularly the range, suggest a non-stationary process? Are the local estimates uniformly significant or restricted to parts of the study area? And, finally, does Monte Carlo simulation suggest spatial variability significantly different from spatial randomness?

A final diagnostic with which to assess additional insight from geographically weighted regression is to examine the local coefficients of determination. Because a model is fitted at each regression point, model summary statistics can also be generated at each point, such as Cook's D, influence statistics, and local adjusted R^2. Depicting local explanatory power, Figure 10 shows that model fit is highest along a belt extending from southeastern to east Africa and continuing northeast to include Egypt and eastern Libya, with values approaching 0.87. Conversely, the model performs least well in coastal central Africa, particularly in Gabon, Equatorial Guinea, coastal Cameroon, and much of Congo (Brazzaville). Minimum values tending to 0.73 would suggest, relative to eastern Africa, that other unspecified determinants are formative drivers of poverty here.

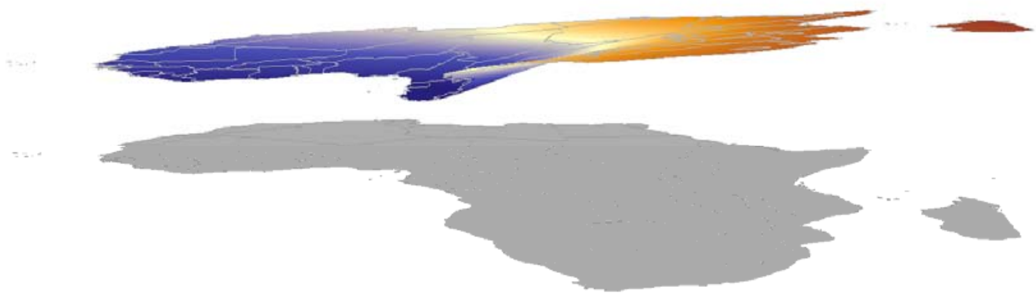

Figure 10: Local Adjusted R2

The foregoing exploration of geographically weighted regression has demonstrated that attending to the local can yield new insights about the processes underlying the construction of poverty. To be sure, social scientists must continue to search for other root causes of poverty and economic stagnation but, we suggest, there remains much to be gleaned from extant theory and data if we recast previous empirical research under the rubric of local forms of analysis, particularly geographically weighted regression.

By focusing on the local construction of poverty we have found just through this partial analysis that some drivers of poverty at the continental level are less effectual in particular regions of the continent and more salient in others with respect to the magnitude of their coefficients. Other drivers are statistically significant in some locales and not significantly different from zero elsewhere. In some cases, directional stability in poverty determinant relationships is non-stationary. We noticed, finally, that our simple poverty model has superior explanatory power for some parts of Africa and performs less well in others. How, though, do we move from theory and analytics to praxis?

Policy Implications and Conclusion

If we understand spatially differentiated drivers of poverty at a sufficiently large geographic scale, then we may simultaneously have strategic and operational cues towards poverty alleviation interventions. While we were unable to procure complete subnational datasets for this analysis we hope to have sufficiently illustrated the added value accruing from a particular form of local analysis—geographically weighted regression. GWR and other techniques of local analysis, informed by rich subnational datasets, can combine to produce highly focused and efficient spatially targeted interventions. We believe that efforts to accelerate poverty reduction and the achievement of the Millennium Development Goals can be profitably informed by an empirical paradigm shift away from the global, cross-country regression framework to one that is sensitive to causal non-stationarity and situates poverty reduction theory and praxis at scale-appropriate geographies.

We would endorse Unwin's (2004 :1519) call—as he argues against the generic one-model-fits-all direct budgetary support mantra currently *en vogue*—that a "focus on local specificity is important since it encourages donors to shift away from essentially uniform economic explanations in order to understand the influence of cultural, social and political factors in shaping the lives of poor people."

However, geographic targeting as a poverty alleviation tool is hardly novel, though its success has been somewhat limited. Past experiments have relied overwhelmingly on developing poverty profiles from rich, small-N surveys and grafting those on to standard census data. Our review of the literature finds that parameter invariance has been implicitly assumed in both poverty modeling and geographic targeting exercises without exception. Measurement and operational geographies have been generally much too large to effect real reductions.

Nonetheless, we see that attending to local sensitivities has already produced some tangible results. UNDP (2005), for example, is beginning to harness and leverage subnational analyses with success in Albania where regional-level reporting is in place to monitor both subnational MDG progress and to monitor more tailored targets against locally specific challenges. The Millennium Project's Millennium Villages initiative further demonstrates how development can be operationally accelerated through place-specific interventions.[16] Focused on local ownership, sustainability, and independence from grid-sourced electricity and water, the project's success stems, to some extent, from scale-appropriate implementation at the community level.

Coupling these kinds of locally tailored, grassroots efforts with robust small area data and local analytic techniques offers an impressive diagnostic and operational tool to combat poverty. Consider how

[16] For a detailed description, see http://www.unmillenniumproject.org/mv/index.htm.

effective programmes and interventions could be if policymakers, aid specialists, and poverty analysts knew what drivers were significantly perpetuating poverty at the provincial or even community level. Perhaps a malaria abatement programme would be most effective in one locale while investments in transport infrastructure would better accelerate poverty reduction in another. Perhaps other interventions such as those focused on literacy and nutrition need to remain global.

In some areas, disaster mitigation investment might effectively reduce poverty since recovering from these kinds of shocks is known to retard development (Collier, 2007). Perhaps, though, in other parts of the continent this kind of spend would be inefficient. Spatially focused efforts can help not only to deliver the most effective intervention for a particular region but also can reduce leakage to the non-poor and increase investment in interventions known to locally covary with levels of poverty. In the effort to achieve equitable distribution, national level policy and interventions may well spread donor aid and investment so thin that it has little efficacy in effecting meaningful change anywhere. Rigorously informed spatial targeting can help to mitigate ineffectual spend.

While theoretically attractive, the operationalization of geographically targeted policy response and intervention remains hampered in much of Africa owing to the dearth of commensurate, reliable, and accessible subnational data. Moving this agenda forward requires that Africa's national statistical agencies, civil society organizations, and international partners invest considerably more towards the continent's statistical infrastructure. Survey design, census cartography, subnational data collection, vetting, and dissemination processes can all benefit from capacity building and productivity investments. There is also, as demonstrated here, an exigency to train a cadre of geoinformation professionals, proficient not just in GIS and remote sensing techniques but also in advanced spatial analytics.

With these kinds of rich subnational data, we believe that a research trajectory around local forms of analysis, especially geographically weighted regression, can pay handsome dividends and can meaningfully contribute to poverty alleviation efforts across the continent, in theoretically understanding the local construction of poverty, in suggesting spatially prescriptive cues on programme design and delivery, and towards maximizing efficient spend.

References

Acemoglu, D. *et al.* (2001a) "Reversal of Fortune: Geography and Institutions in the Making of the Modern World Income Distribution," National Bureau of Economic Research Working Paper 8460. Cambridge, MA: NBER; http://www.nber.org/papers/w8460.

_____ (2001b) "The Colonial Origins of Comparative Development: An Empirical Investigation," *The American Economic Review*, **91**(5): 1369-1401.

Appleton, D.R. *et al.* (1996) "Ignoring a Covariate: An Example of Simpson's Paradox," *The American Statistician*, **50**(4): 340-41.

Barro, Robert and Jong-Wha Lee (1994) "Losers and Winners in Economic Growth," in the *Proceedings of the Annual World Bank Conference on Development Economics*, *pp.* 267-97.

Bigman, D. and H. Fofack (2000) "Geographical Targeting for Poverty Alleviation: An Introduction to the Special Issue," *The World Bank Economic Review*, **14**(1): 129-145.

Bloom, David and Jeffrey Sachs (1998) "Geography, Demography, and Economic Growth in Africa," *Brookings Papers on Economic Activity*, **1998**(2): 207-295.

Bongaarts, John (1994) "Population Policy Options in the Developing World," *Science*, **263**: 771-76.

Collier, Paul (1998) "The Political Economy of Ethnicity," in B. Pleskovic and J. Stiglitz, *eds.*, *Proceedings of the Annual World Bank Conference on Development Economics*.

Collier, Paul (2007) *The Bottom Billion: Why the Poorest Countries Are Failing and What Can Be Done About It.* Oxford, UK: Oxford UP.

Collier, Paul and Jan Willem Gunning (1999) "Explaining African Economic Performance," *Journal of Economic Literature*, **37**(1): 64-111.

Coulombe, Harold and Andrew McKay (1996) "Modeling Determinants of Poverty in Mauritania," *World Development*, **24**(6): 1015-1031.

Diamond, J. (1997) *Guns, Germs, and Steel: The Fate of Human Societies*. New York, NY: Norton & Co.

Easterly, W. and N. Levine (1997) "Africa's Growth Tragedy: Policies and Ethnic Division," *Quarterly Journal of Economics*, **112**(4): 1203-1250.

Fotheringham, A.S. *et al.* (2002) *Geographically Weighted Regression: The Analysis of Spatially Varying Relationships*. West Sussex, UK: Wiley.

Glewwe, P. (1991) "Investigating the Determinants of Household Welfare in Cote d'Ivoire," *Journal of Development Economics*, **35**(2): 307-337.

Hentschel, J. *et al.* (1998) "Combining Census and Survey Data to Study Spatial Dimensions of Poverty," Policy Research Working Paper 1928. Washington, DC: World Bank.

Hibbs, Douglas *et al.* (2004) "Geography, Biogeography, and Why Some Countries Are Rich and Others Are Poor," *Proceedings of the National Academy of Sciences of the United States of America*, **101**(10): 3715-20.

Knapp, Thomas (1985) "Instances of Simpson's Paradox," *The College Mathematics Journal*, **16**(3): 209-11.

Kyereme, S.S. and E. Thorbecke (1991) "Factors Affecting Food Poverty in Ghana," *Journal of Development Studies*, **28**(1): 38-52.

Maasoumi, Esfandiar and Maria Ana Lugo (2008) "The Information Basis of Multivariate Poverty Assessments," *pp.* 1- 29 in Nanak Kakwani and Jacques Silber, *eds.*, *Quantitative Approaches to Multidimensional Poverty Measurement.* Hampshire, UK: Palgrave Macmillan.

McKinley, Terry (2006) "What Is Poverty? Good Question," *IPC One Pager*, **26**. Brasilia, Brazil: International Poverty Centre.

Ravaillion, Martin (1996) "Issues in Measuring and Modelling Poverty," *The Economic Journal*, **106**(438): 1328-43.

Rodriquez, Francisco (2007) "Policymakers Beware: The Use and Misuse of Regressions in Explaining Economic Growth," *IPC Policy Research Brief*, **5**. Brasilia, Brazil: International Poverty Centre.

Sachs, J. (2001) "Tropical Underdevelopment," National Bureau of Economic Research Working Paper 8119. Cambridge, MA: NBER.

_____ (2003a) "Institutions Matter, but Not for Everything," *Finance and Development*, June 2003.

_____ (2003b) "Institutions Don't Rule: Direct Effects of Geography on Per Capita Income," National Bureau of Economic Research Working Paper 9490. Cambridge, MA: NBER.

Sachs, J. and A. Warner (1997) "Sources of Slow Growth in African Economies," *Journal of African Economics*, **6**: 335-76.

Sahn, David and David Stifel (2002) "Robust Comparisons of Malnutrition in Developing Countries," *American Journal of Agricultural Economics*, **84**(3): 716-35.

Sen, A. (1985) *Commodities and Capabilities*. Amsterdam: North-Holland.

_____ (1987) *The Standard of Living*. Cambridge, MA: Cambridge UP.

Serieux, John and Terry McKinley (2008) "Is Financial Liberalization a Flop? An Africa Assessment," *IPC One Pager*, **48**. Brasilia, Brazil: International Poverty Centre.

Simpson, E.H. (1951) "The Interpretation of Interaction in Contingency Tables," *Journal of the Royal Statistical Society, Series B (Methodological)*, **13**(2): 238-41.

Tobler, Waldo (1970) "A Computer Model Simulation of Urban Growth in the Detroit Region," *Economic Geography*, **46**(2): 234-40.

UNDP (2005) "Advancing Development Through Sub-national MDG Reports: A Case Study of Albania," UNDP Good Practices Series. Accessed February 15, 2008 at http://www.undp.org/mdg/goodpractices.html.

Unwin, Tim (2004) "Beyond Budgetary Support: Pro-Poor Development Agendas for Africa," *Third World Quarterly*, **25**(8): 1501-23.

Van de Walle, N. (2001) *African Economies and the Politics of Permanent Crisis, 1979-1999*. Cambridge, UK: Cambridge UP.

Wagner, Clifford (1982) “Simpson’s Paradox in Real Life,” *The American Statistician*, **36**(1): 46-48.

Woods, Dwayne (2004) “Latitude or Rectitude: Geographical or Institutional Determinants of Development,” *Third World Quarterly*, **25**(8): 1401-14.