Development of High-Resolution Gridded Poverty Surfaces
Bill and Melinda Gates Foundation Contract #21989

Final Report: Development of High-Resolution Gridded Poverty Surfaces

Jan 10th 2014

Dr Andrew Tatem (Co-PI)
Dr Peter Gething (Co-PI)
Dr Carla Pezzulo
Dr Dan Weiss
Dr Samir Bhatt

UNIVERSITY OF Southampton

UNIVERSITY OF OXFORD

1.0 OVERVIEW

Improved understanding of geographic variation and inequity in health status, wealth, and access to resources within countries is increasingly recognized as central to meeting development goals. Development indicators assessed at national scales can often conceal important inequities, with the rural poor often least well represented. As international funding for development comes under pressure, the ability to target limited resources to underserved groups becomes crucial. Monitoring inequalities for targeting interventions requires a reliable and detailed evidence base. While high-resolution spatial data on population distributions in resource poor areas are now becoming available (e.g. www.worldpop.org.uk), comprehensive information on demographic, health and wealth attributes of those populations remain only usable at highly aggregated regional levels through national household surveys [1].

The Demographic and Health Survey (DHS) and Living Standards Measurment Survey (LSMS) programs have been leaders in collecting and providing cluster-randomised survey data on core development indicators. In addition to their standard open-source data files in which survey results are tabulated by first-order sub-national regions (for example at province or state level) and urban/rural strata, more recent surveys now provide geocoded data for individual clusters. The availability of the GPS coordinates for DHS (and Malaria Indicator Surveys (MIS) and Aids Indicator Surveys (AIS)) and LSMS clusters provides, for the first time, highly resolved locational information that can be linked with survey outputs for quantifying demographic, health and economic status heterogeneities and inequities.

Here we present pilot outputs of a novel spatial statistical methodology for the production of gridded surfaces of DHS and LSMS-derived variables, focusing on poverty mapping in six countries: Kenya, Uganda, Tanzania, Malawi, Nigeria and Pakistan. A Bayesian geostatistical modeling framework, following approaches constructed for the Malaria Atlas Project [2], has been established to exploit spatiotemporal relationships within the data, leverage ancillary information from an extensive set of covariates, and rigorously handle uncertainties at all stages to generate robust output surfaces with accompanying confidence intervals. We demonstrate the application of the approach to mapping both asset-based poverty metrics from DHS data (for Kenya, Uganda, Tanzania and Pakistan) and consumption-based poverty metrics from LSMS data (for Malawi, Nigeria, Uganda and Tanzania).

2.0 DEFINING METRICS OF POVERTY

While there is wide consensus on the need to reduce poverty-related inequalities in health, there is equally wide variation in the definition of poverty itself. Conventionally, the most common unidimensional approach to measuring poverty is the monetary approach where poverty is related to income or to a money metric measure such as expenditures. As described in Alkire and Foster [3] the monetary approach considers as ‘poor’ those who have a shortfall in consumption or income or whose income stands below a set poverty line. ‘Unidimensional methods can be applied when a well-defined single-dimensional resource variable, such as income, has been selected as the basis for poverty evaluation’ [3]. The monetary definition of poverty still remains central in many approaches today. Millennium Development Goal 1, aimed at eradicate extreme poverty and hunger, implicitly assumes wellbeing being determined by income, given its attempt 'to reduce the proportion of the population living on less than US$1 a day by one half between 1990 and 2015' [4]. However, more recent poverty investigation has taken into account a broader variety of social indicators and assets, together with income and consumption. Here, we describe the application of our mapping approaches to both asset-based multidimensional measures of poverty (as described in section 2.1) and consumption-based monetary measures of poverty (as described in section 2.2).

2.1 ASSET BASED POVERTY MEASURES

Recent poverty analyses have seen a trend towards multidimensional metrics, due to the quality (regularity and comparability) of income/expenditure data being often poor in many developing countries, especially Sub-Saharan African ones, which are generally regarded as those showing the most poverty and extreme poverty. Moreover, well-being and poverty are now often seen as multidimensional phenomena, with the well-being of an individual depending not only on income, but also on several other dimensions or capabilities such as health, education and empowerment, amongst others.

The DHS program provides data on core development indicators that are available at household level. GPS coordinates from DHS are available at cluster level, resulting in highly detailed locational information. Therefore, using the 'Multidimensional Poverty Index' (MPI) [5] as a poverty metric allowed us here to maintain the information at a disaggregated level, reducing the loss of detail. Moreover, our intention was to capture as many different aspects of poverty as possible, from the assets and utilities available in the household to the level of education of its members and health.

2.1.1 The Multidimensional Poverty Index (MPI)

The Multidimensional Poverty Index (MPI) was proposed by Alkire and Foster in 2007[5] –and then further described in Alkire and Santos [6]-as a result of a joint effort of the Oxford Poverty and Human Development Initiative (OPHI) of Oxford University and the Human Development Report Office of the United Nations Development Programme (UNDP). The rationale for the construction of the MPI shares the same idea of the Human Development Index, being a score composed by the same three dimensions, health, education and standard of living (Figure 1). The MPI merges the money-based and utility measures with health and education elements, which attempt to function also as proxies of availability and accessibility to services.

The MPI measures deprivation instead of possession and *"shows the number of people who are multidimensionally poor (suffering deprivations in 33% of weighted indicators) and the number* of deprivations with which poor households typically contend"[7]. The three aforementioned dimensions composing the MPI consist of a total of ten indicators (Figure 1). The health dimension includes nutrition and child mortality indicators. The education dimension is composed of child enrolment and years of schooling and the standard of living dimension lists several assets that might be owned by every household.

The DHS questionnaires typically report all the necessary information needed to create the MPI, with some exceptions for older surveys. Using the DHS microdata, the index can be constructed by different population subgroups at household or cluster level as well as at higher levels (Region, Country). It can be also decomposed by dimension to show how the structure of poverty differs between different groups.

Following the methodology proposed by Alkire and Santos [6] the process of assembling the index starts with assigning to each household a weighted score following each indicator threshold, in order to define the household as poor or not poor. Specifically, the health thresholds are having at least one household member who is malnourished and having had one or more children die. The education thresholds are having no household member who has completed five years of schooling and having at least one school-age child (up to grade 8) who is not attending school. The standard of living thresholds relate to not having electricity, not having access to clean drinking water, not having access to adequate sanitation, using “dirty” cooking fuel (dung, wood or charcoal), having a home with a dirt floor, and owning no car, truck or similar motorized vehicle, and owning, at most, one of these assets: bicycle, motorcycle, radio, refrigerator, telephone or television.

To identify the multidimensionally poor, the deprivation scores for each household are summed to obtain the household deprivation. A cut-off (1/3 of the indicators), is used to distinguish between the poor and non-poor. If the deprivation score is 3 or greater, that household (and everyone in it) is multidimensionally poor. An additional step is required to calculate an aggregate measure of the MPI. The MPI value is the product of two measures: the **headcount ratio** and the **intensity of poverty**. The headcount ratio is the proportion of the population who are multidimensionally poor. In our work, this ratio has been calculated at the cluster level as the number of people who are multidimensionally poor on the total population, by cluster.

The intensity of poverty reflects the proportion of the weighted component indicators in which, on average, poor people are deprived. For poor households only, the deprivation scores are summed and divided by the total number of indicators and by the total number of poor persons. The MPI, product of headcount ratio and intensity of poverty, ranges between 0 and 1 where 0 is multidimensionally not deprived and 1 is multidimensionally deprived. Headcount ratio, intensity of poverty as well as the total MPI value can be aggregated at cluster level, as well as at national level.

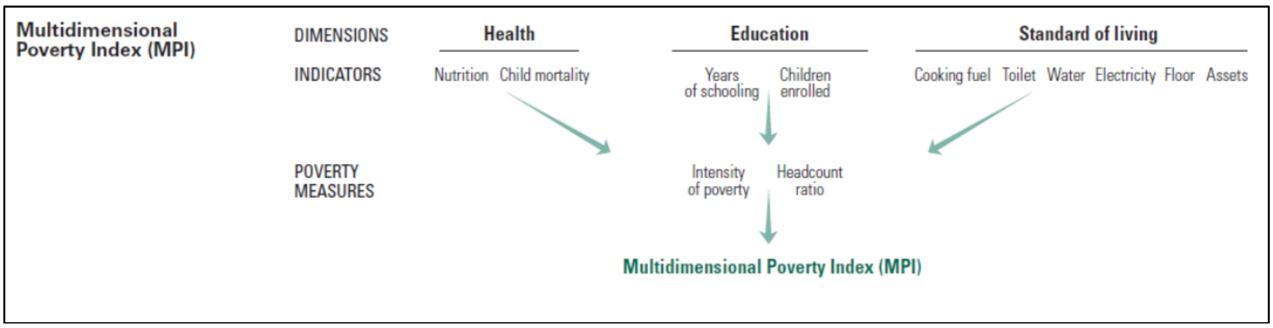

Figure 1. The composition of the Multidimensional Poverty Index (MPI). Source: [6]

2.1.2 Calculating the MPI for Kenya, Uganda, Tanzania and Pakistan

The most recent DHS datasets were obtained for four study countries for which MPI mapping was undertaken here, and the features of each dataset are outlined in table 1. For Tanzania, two recent datasets were obtained and pooled to provide a greater coverage spatially. The MPI construction methodology described above in section 2.1.1 was applied to the DHS datasets. For the construction of the MPI metrics, there were cases where the survey did not collect part of the information needed to construct the MPI. This occurred in the case of the Tanzania 2012 AIS and the Pakistan 2006/7 DHS, and the following adjustments were made to account for these:

<u>Tanzania AIS 2012</u>

- NUTRITION: Neither women or children nutrition information were reported, therefore the nutrition weight was added to the child death indicator to create the Health Dimension.
- CHILD DEATH: the available information on Child Death referred only to those children born during the 6 years before the survey, while Standard DHS Survey usually collects this information for all children ever born. The child death information was used without modification in this case.

<u>Pakistan DHS 2006/7</u>

- NUTRITION: Neither women or children nutrition information were reported, therefore the nutrition weight was added to the child death indicator, to create the Health Dimension.

Figure 2 shows the locations of the survey clusters for East Africa and Pakistan and the values of the MPI for each one. In the remainder of this study for MPI we model and map the three East African countries together as a regional block. Future work should ideally examine the benefits, drawbacks and sensitivities of outputs to undertaking this, rather than treating each country separately.

Country	Year	Name of survey and Executing Agency	Universe	Number of:					Sampling method and source
				Women interviewed[1]	Households interviewed	Household members[2]	Clusters (Primary Sampling Unit – DHS survey)	Final number of clusters[3]	
Kenya									
	2008	Kenya Demographic and Health Survey 2008-2009 - Kenya National Bureau of Statistics	All women 15-49	8,444	9,057 74% rural 26% urban	38,019 49 % male 51% female	400	397 68% rural 32% urban	Two-stage sample based on the 1999 Population and Housing Census
Tanzania									
	2011/12	2011-12 Tanzania HIV/AIDS and Malaria Indicator Survey - National Bureau of Statistics (Census Office)	All women and all men 15-49	10,967	10,040 Mainland: 8,727 78% rural 22% urban Zanzibar: 1,313	54,020 48% male 52% female	583	573 77% rural 23% urban	Two-stage sample based on the 2002 Population and Housing Census (PHC)
	2010	Tanzania Demographic and Health Survey 2010 - Tanzania	All women 15-49	10,139	9,623 Mainland: 9,377 74% rural 26% urban	50,414 49% male 51% female	475	458 78% rural 22% urban	Two-stage sample based on the 2002 Population and Housing

[1] Data from individual questionnaire on women

[2] Data consistent with the Reports, indicating the respective numbers of successful interviews. Not weighted. The final datasets used in this project may have different numbers given the due adjustments (exclusion of missing cases as for the MPI methodology, exclusion of clusters with missing GPS information)

[3] Net of adjustments (excluding non-geo-referenced clusters from GPS dataset – as MIS source of reference - and other missing cases, e.g. after calculating poverty metrics.)

		National Bureau of Statistics			Zanzibar: 246				Census (PHC)
Uganda									
	2011	Uganda Demographic and Health Survey 2011 -Uganda Bureau of Statistics Kampala, Uganda	All women age 15-49	8,674	9,033 81% rural 19% urban	44,250 49% male 51% female	404	400 71% rural 29% urban	Two stages sample based on Uganda National Household Survey and 2002 Population Census sample frame
Pakistan									
	2006/7	Pakistan Demographic and Health Survey 200 6-07 National Institute of Population Studies Islamabad, Pakistan	Ever married women 15-49	10,023[4]	95,441 (Total) 60% rural 40% urban 86,186 (Short Household Questionnaire) 9,255 (subsample for women selection -Long Household Questionnaire- 10% subsample)	688,937 49% female 51% male	972	957 60% rural 40% urban	Two-stage, stratified, random sample design based on the 1998 population census.[5]

Table 1. The features of the geolocated household survey datasets used in construction of the Multidimensional Poverty Indices (MPI). Source: [6]

[4] Number of women interviewed using the Short Questionnaire. The MPI measures are derived from the Short Questionnaire sample

[5] The sample for the 2006-07 PDHS represents the population of Pakistan excluding the Federally Administered Northern Areas (FANA) and restricted military and protected areas. Although the Federally Administered Tribal Areas (FATA) were initially included in the sample, due to security and political reasons, it was not possible to cover any of the sample points in the FATA.

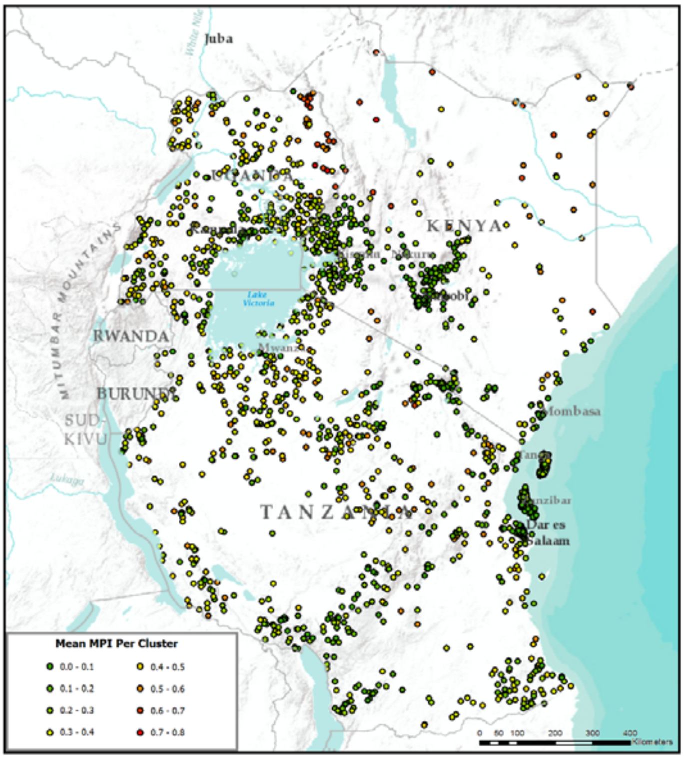

(a)

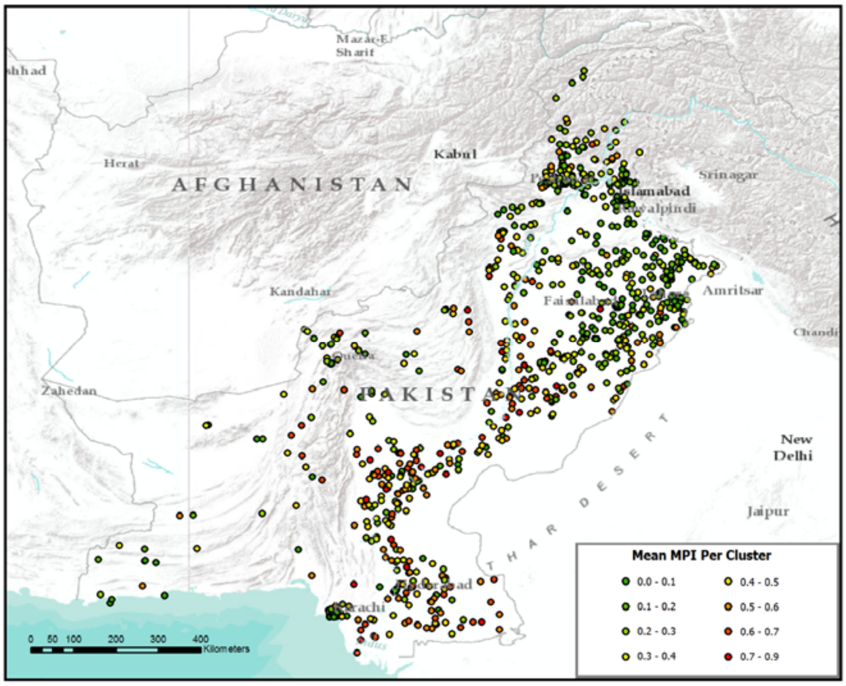

(b)

Figure 2. Input survey clusters for (a) East Africa, displaying the cluster-level MPI headcount ratio derived from the Kenya, Uganda and Tanzania surveys outlined in table 2, and (b) Pakistan displaying the cluster-level MPI headcount ratio derived from the 2006/7 DHS survey.

2.2 CONSUMPTION-BASED WELFARE INDICATORS AND POVERTY MEASURES

The household surveys used in this project for creating consumption-based poverty measures are part of the LSMS programme held by the World Bank together with Countries' Bureaus of Statistics. The LSMS aims at improving the quality of household data collected by statistical offices in developing countries in order to refine the measures of households' standards of living and poverty status. The LSMS surveys collect information on total household expenditures on an item-by-item basis, including expenditure on food, beverages and tobacco, non-durable goods and frequently purchased services; semi-durable and durable goods and services; and non-consumption expenditures [8]. The World Bank, among other money metric indicators (e.g. income based measures), use a consumption based metric as an indicator for the estimation of wellbeing. There are several steps to be followed for estimating a poverty measure using the household level data:

"Poverty analysis requires three main elements. The first component is a welfare indicator to rank all the population accordingly. The second element is an appropriate poverty line to be compared against the chosen indicator in order to classify individuals into poor and non-poor. The final component is a set of measures that combine individual welfare indicators into an aggregate poverty figure."[9]

1) The Welfare Indicator

When poverty is estimated using monetary measures, income and consumption are generally the two most common measures used. Most analysts [10] nowadays claim consumption to be a better measure for poverty, in the family of quantitative metrics. As argued by Coudouel et al. [11] consumption is a direct indicator of a family or individual outcomes, while income is not always directly related to the achievement of the basic needs. Moreover, as it is also argued, consumption data are usually more reliable than income data. The reason why expenditure as proxy of consumption is preferred to income measures, in the framework of monetary measures of poverty, is largely discussed in the literature [11,12,13]. As summarized by the World Bank – Measuring Poverty Section[6] and in a World Bank report [14], consumption captures better than income current basic needs of a household, therefore it is closely related to a person's well-being; especially in developing countries where the rate of informal economy is high, consumption data are often better collected and measured than income. Moreover, consumption may better reflect a household's ability to meet basic needs. Therefore generally, at both national and international levels, consumption is nowadays a more desired metric for poverty analysis.

LSMS generally report poverty measures based on a consumption metric as an indicator for the estimation of wellbeing among other money metric indicators (e.g. income based measures). Therefore all their estimates on poverty are mostly constructed on family expenditures, functioning as base indicators for applying different poverty lines.

The **Construction of the Consumption Aggregate**: a few general principles are adopted when constructing the consumption aggregate and these are generally common to every survey used in this project. In this section we will only report basic information about the methodology for

[6] http://go.worldbank.org/W3HL5GD710

creating the consumption aggregate. Details can be found on individual country reports released with the relevant survey.

Some general methods are reported here following the Appendix A methodology for consumption of the Tanzania NPS Report for 2010-11 and the “Guidelines for constructing consumption aggregates for welfare analysis” by Deaton and Zaidi [10], where they propose the steps for constructing a welfare indicator based on consumption expenditures and describe the necessary adjustments for aggregating data. Their methodology is usually applied, with adjustment, to several living standards measurement study surveys – (LSMS). The theoretical basis for the construction of the consumption aggregate proposed by Deaton and Zaidi [10] is the money metric utility as described by Samuelson [15]. In their model, they approximate a money metric utility by adding up all the household's expenditures, and dividing by a Paasche index of prices, in order to adjust for price differences (different prices in different countries for comparable goods).

The process of individualization of the aggregation of a consumption indicator is generally complex. Some adjustments might be necessary to ensure that the aggregation process leads to the desired measures. The first rule when aggregating a consumption indicator is to include all the possible sources, without excluding components that contribute to the family welfare. Second, generally both market and non-market transactions are to be included. Third, given that expenditure is not equivalent to consumption, corrections are needed when including housing and durable goods, assuming that not all these goods are consumed. Fourth, a common reference period should be chosen. Different items can be reported per different periods of time; however the final consumption aggregate is generally reported per 28 days or at annual level.

The different components included in the consumption aggregate are generally the food component (food purchased in the market, food eaten away from home, food produced at home, received as gift), accounting only for those food actually consumed, measured it unit values. Non-food component includes data on an extensive range of non-food items, such as household utilities, health expenses and education. A few decisions on what to include and what to exclude among some services (such as debts or remittances) have to be done. In [10] a detailed discussion on the rationale for including health expenditures and education expenditures in the calculation of consumption of non-food items is reported. They propose that the decision on whether or not to include these expenses should depend on the elasticity of the health and education expenditures on total expenditure. In the case of high elasticity these expenses have to be included. Moreover, ownership of durable goods is included in the component with few adjustments and considering the quality and quantity of data actually available in the survey. Additionally, items indicating housing conditions and ownership have to be considered, accounting for the difficulties of estimating elements like rents, for example, and choosing according to the different situations.

All of the expenditures for utilities and amenities reported by the households have to be included in the calculation. Deaton and Zaidi [10] then propose some solutions to adjust for cost of living differences and for household size and composition. Generally temporal and spatial price adjustments are made to adjust consumption to real terms, by applying a spatial and time deflator. Deaton and Zaidi propose the use of the Paasche price index, which assigns a weight to each household. However, in the case of Tanzania NPS, the Fisher price indices based only on food items were employed. The last adjustment then concerns the household composition: in this step, the welfare indicator goes from a measure of standard of living defined at the household level to another at the individual level. In order to take into account the differences in needs between households and intra-household inequalities the equivalence scales are usually used for comparing consumption aggregates of households with different demographic compositions. In particular, Deaton and Zaidi apply the *arbitrary method* to derive the equivalence scales.[7]

In the case of our poverty mapping here, given the use of the Intenational Poverty Lines, which by definition are constructed at a *per capita* level, the household consumption measure is divided by the number of household members. This choice was made following the PovcalNet procedures: *"The per capita income/consumption used in PovcalNet is household income/consumption expenditure dividing by the household size"*[8]

2) <u>The choice of the Poverty Line</u>
This brings us to the second step for measuring poverty, which concerns the choice of the poverty line. After the selection of the indicator - commonly based on consumption in the case of adopting the monetary approach - and after having identified the unit of analysis, a poverty line is adopted. As introduced, we apply the International absolute poverty lines converted on a monthly base (or annual base, depending on the consumption indicator base) and accounting for current prices.

These thresholds that identify who is poor and who is not poor, were set at the international absolute poverty lines of $1.25 a day per person (2005 Purchasing Power Parity) and $2 a day per person (2005 Purchasing Power Parity), following the updated estimations done for measuring the most famous 'dollar a day' [17]. The Purchasing Power Parity (PPP) conversion factors are mostly used for international comparisons of prices. Therefore, the PPP values used to convert the dollar a day to local currencies take into account the differences in purchase power for the same goods in different countries. The corresponding values in country currency prices are derived through the World Bank PovCalNet website.[9] The relevant Consumer Price Index (corresponding with the year of the survey under analysis) is then applied following, among others, the methodology available on USAID Poverty Tools website in order to adjust for the corresponding prices and inflation.[10] The relevant CPI is taken from the PovcalNet website.[11]

---

[7] The equivalence scale is an index converting nominal incomes of heterogeneous households in comparable measures of well-being. This index can be interpreted as the differential cost of having a given household size and composition with respect to a «benchmark» household type. Equivalence scales are the traditional measure proposed by the OECD. For more detailed information on equivalence scales see 16. Bellu' LG, Liberati P (2005) Equivalence Scales General Aspects. Food and Agriculture Organization of the United Nations, FAO.

[8] http://iresearch.worldbank.org/PovcalNet/index.htm?0,2

[9] PovcalNet: the on-line tool for poverty measurement developed by the Development Research Group of the World Bank. PovcalNet - an online poverty analysis tool

[10] "Calculating PPP Conversion Factors and "$1-a-day" Poverty Lines" Adapted by Don Sillers of USAID, from the Annex to National and International Poverty Lines: An Overview, available at Poverty Assessment Tools

[11] http://iresearch.worldbank.org/PovcalNet/index.htm?4 (then go at the bottom of the page and click on Uganda 2009 – CPI and Population table)

3) Aggregate Poverty Metric
As a last step, following Coudouel et al.'s [11] discussion on poverty measurements and analysis, it is necessary to choose and estimate a poverty measure. This measure is a statistical function that translates the relation between the indicator and the poverty line into one aggregate number. The measures usually refer to the entire population or to subgroups and can be calculated on a household basis or on an individual basis.

The measures most commonly used can indicate the incidence of poverty (as in the case of the headcount ratio), the depth of poverty (as in the case of the poverty gap index), or the severity of poverty (as in the case of the squared poverty gap). Here, we compare the consumption aggregate per capita against the poverty line and assign the status of poor or not poor to each member of the family whose consumption falls below the poverty line. We then calculate the headcount ratio per person for each Enumeration Area Unit (or Primary Sampling Unit) dividing the number of those members considered poor by the total number of members in the EA. We therefore link each EA Latitude and EA Longitude with a Unique EA Identifier.

2.2.1 Calculating Consumption-based metrics for Uganda, Nigeria, Tanzania and Malawi

The data and methods used to calculate the consumption-based poverty measures from national household survey data for each of the four countries for which consumption-based poverty maps were produced are described below in table 2.

COUNTRY and YEAR	
	Uganda 2009/10
SURVEY TYPE and SOURCE	
	Sample survey data: Uganda National Panel Survey (UNPS) 2009/10 Principal Investigator: Uganda Bureau of Statistics (UBOS) Other producers: Government of Netherlands and World Bank Living Standards Measurement Study Study Type: Living Standards Measurement Study (World Bank) http://microdata.worldbank.org/index.php/catalog/1001/study-description
SURVEY SAMPLE	
	Starting in 2009/10, the UNPS has been set out to track and reinterview 3,123 households that were distributed over 322 enumeration areas (EAs), selected out of the 783 EAs that had been visited by the Uganda National Household Survey (UNHS) in 2005/06. The UNPS EAs covered all 34 EAs visited by the UNHS 2005/06 in Kampala District, and 72 EAs (58 rural and 14 urban) in each of the (i) Central Region with the exception of Kampala District, (ii) Eastern Region, (iii) Western Region, and (iv) Northern Region.
POVERTY LINES	
	The dollar a day poverty lines were converted in local currency (UGS shillings) using the value of 2005 Purchasing Power Parity as published on PovCal Net web site [18]. The 2005 PPP for Uganda is equivalent to 744.618 UGS. An update at current prices was then performed using the

	Consumer Price Index for 2009/10 divided by CPI 2005 (1.442) available on the World Bank Website, on the Povcal Net page [18] or on the Indicators page [19] The consumer Price Index is applied following the methodology available on USAID Poverty Tools website in order to adjust for the 2009/10 prices and inflation [20]

COUNTRY and YEAR	
	Nigeria 2010/11
SURVEY TYPE and SOURCE	
	Sample survey data: Nigeria General Household Survey–Panel 2010/11 Principal Investigator: National Bureau of Statistics Other producers: The World Bank, Federal Republic of Nigeria Study Type: Living Standards Measurement Study (World Bank) http://microdata.worldbank.org/index.php/catalog/1002/study-description
SURVEY SAMPLE	
	The sample frame includes all thirty-six (36) states of the federation and Federal Capital Territory (FCT), Abuja. Both urban and rural areas were covered and in all, 500 clusters/EAs were canvassed and 5,000 households were interviewed. These samples were proportionally selected in the states such that different states have different samples.
CONSUMPTION DATA & METHOD (specific notes)	
	Citing the Nigeria GHS Basic Information Document, Section 8 on the Calculation of consumption aggregate: *"The consumption aggregates are computed from the expenditure sections of the questionnaire for general food and non-food expenditures. In addition to this, educational expenditures are obtained from the education of the questionnaire for both post-planting and post-harvest. In the case of the post-harvest visit, a housing expense section was included in the questionnaire and these data were used in the computation of the consumption aggregate. A housing expense section was not included in the post-planting questionnaire."* Therefore, in our analysis we used the provided consumption aggregate based on the post-harvest questionnaire. Appendix 3 of the same document also reports the methodology of calculation and lists the items included in the aggregate. The Nigeria GHS household consumptions were annualized, therefore we applied the international absolute poverty lines on annual base and accounting for the 2010/11 prices.
POVERTY LINES	
	The dollar a day poverty lines were converted in local currency (NGN nair) using the value of 2005 Purchasing Power Parity as published on PovCal Net web site [18]. The 2005 PPP for Nigeria is equivalent to 78.583 NGN. An update at current prices was then performed using the Consumer Price Index for 2010 divided by CPI 2005 (1.709) available on the World Bank Website, on the Povcal Net page [18] or on the Indicators page [19] The consumer Price Index is applied following the methodology available on USAID Poverty Tools website in order to

	adjust for the 2010 prices and inflation [20]
NOTES ON GPS	
	The GPS coordinates for 2010/11 Nigeria GHS are directly provided by the LSMS Team (World Bank). The documentation provided together with the GPS coordinates clearly states that these locations are a preliminary version not for distribution outside the team. The coordinates will be revised and the official version will be released to the public in the next months.

COUNTRY and YEAR	
	Tanzania 2010
SURVEY TYPE and SOURCE	
	Sample survey data: Tanzania National Panel Survey Wave 2 – 2010-11 Principal Investigator: National Bureau of Statistics and World Bank Study Type: Living Standards Measurement Study (World Bank) http://microdata.worldbank.org/index.php/catalog/1050/study-description
SURVEY SAMPLE	
	The sample design for the second round of the NPS revisits all the households interviewed in the first round of the panel, as well as tracking adult split-off household members. The original sample size of 3,265 households was designed to representative at the national, urban/rural, and major agro-ecological zones. The total sample size was 3,265 households in 409 Enumeration Areas (2,063 households in rural areas and 1,202 urban areas). The total sample size for the second round of the NPS has a total sample size of 3924 households. This represents 3168 round-one households, a re-interview rate of over 97 percent. In addition, of the 10,420 eligible adults (over age 15 in 2010), 9,338 were re-interviewed, a re-interview rate of approximately 90 percent.
CONSUMPTION DATA & METHOD (specific notes)	
	Peculiarities of Tanzania NPS 2010-11 *"It should be noted that although poverty analysis based on the NPS uses the same methodology as the Household Budget Surveys (HBS), the findings in the NPS are not directly comparable to those of the HBS. This is largely attributed to the different technique of collecting consumption data in the two surveys. […] therefore HBS will remain to be the official source of the incidence of poverty in the country."* Tanzania National Panel Survey Report - Wave 2, 2010/11
POVERTY LINES	
	The dollar a day poverty lines were converted in local currency (TZS shillings) using the value of 2005 Purchasing Power Parity as published on PovCal Net web site [18]. The 2005 PPP for Tanzania is equivalent to 482.451 TZS. An update at current prices was then performed using the Consumer Price Index for 2010 divided by CPI 2005 (1.508) available on the World Bank Website, on the Povcal Net page [18] or on the Indicators page [21] The consumer Price Index is applied following the methodology available on USAID Poverty Tools website

in order to adjust for the 2010 prices and inflation [20]

COUNTRY and YEAR	
	Malawi 2010-2011
SURVEY TYPE and SOURCE	
	Sample survey data: Third Integrated Household Survey (IHS3) Primary Investigator: National Statistical Office (NSO) - Ministry of Economic Planning and Development (MoEPD) Other producer: World Bank Study Type: Living Standards Measurement Study (World Bank) URL:http://microdata.worldbank.org/index.php/catalog/1003/study-description#page=overview&tab=study-desc
SURVEY SAMPLE	
	The final sample includes 12,271 households and 768 EAs (Enumeration Area Units or Primary Sampling Units). A stratified two-stage sample design was used for the IHS3. For further details on sample design and sample size see the sampling procedure section on Third Integrated Household Survey (IHS3) Report [22].
POVERTY LINE	
	The dollar a day poverty lines were converted in local currency (Malawian Kwacha) using the value of 2005 Purchasing Power Parity as published on PovCal Net web site [18] The 2005 PPP for Malawi is equivalent to 56.922 MWK. An update at current prices was then performed using the Consumer Price Index for 2011 divided by CPI 2005 (1.676) available on the World Bank Website, on the Povcal Net page [18] or on the Indicators page [21] The Consumer Price Index is applied following the methodology available on USAID Poverty Tools website in order to adjust for the 2010 prices and inflation [20]

Figures 3,4,5 and 6 below show the input survey cluster locations for each of the four countries mapped, with the cluster points coloured by the poverty headcounts for <$1.25 a day and <$2 a day thresholds.

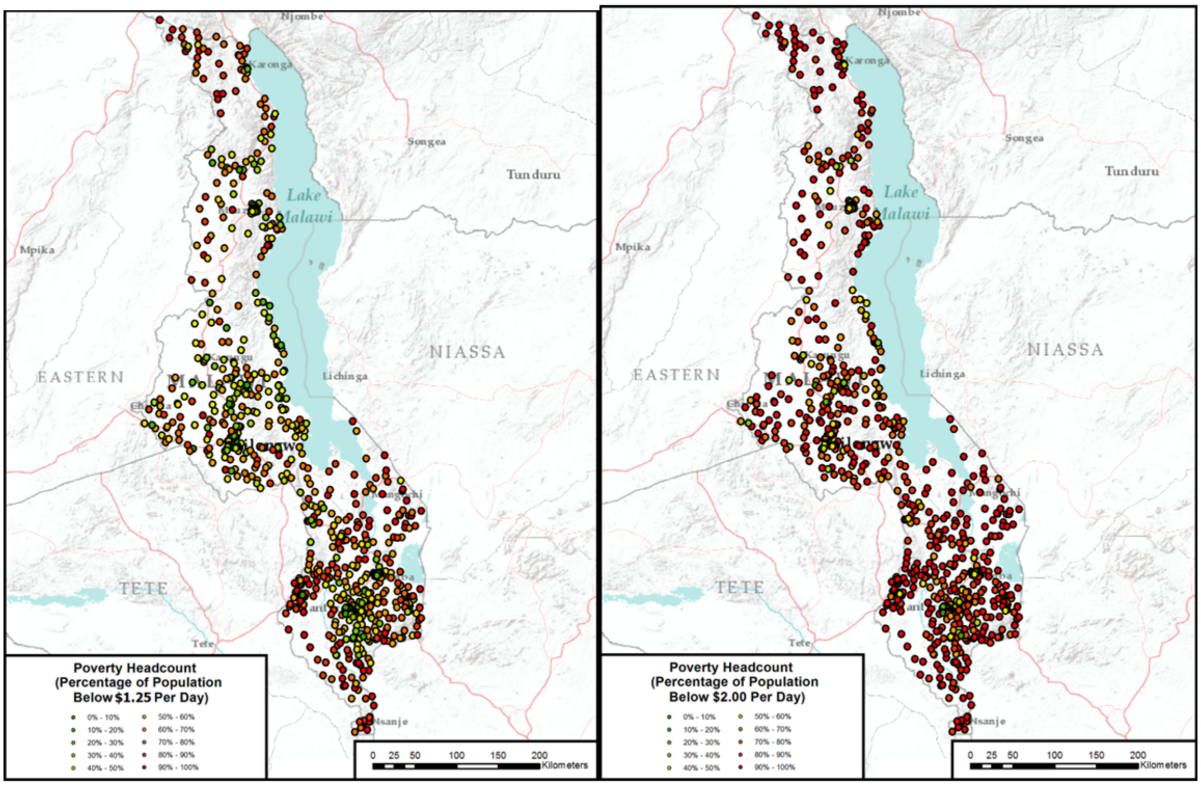

(a) (b)

Figure 3. Input survey clusters for Malawi displaying the cluster-level consumption-based poverty headcount ratio derived from household surveys, for (a) below $1.25 a day and (b) below $2 a day.

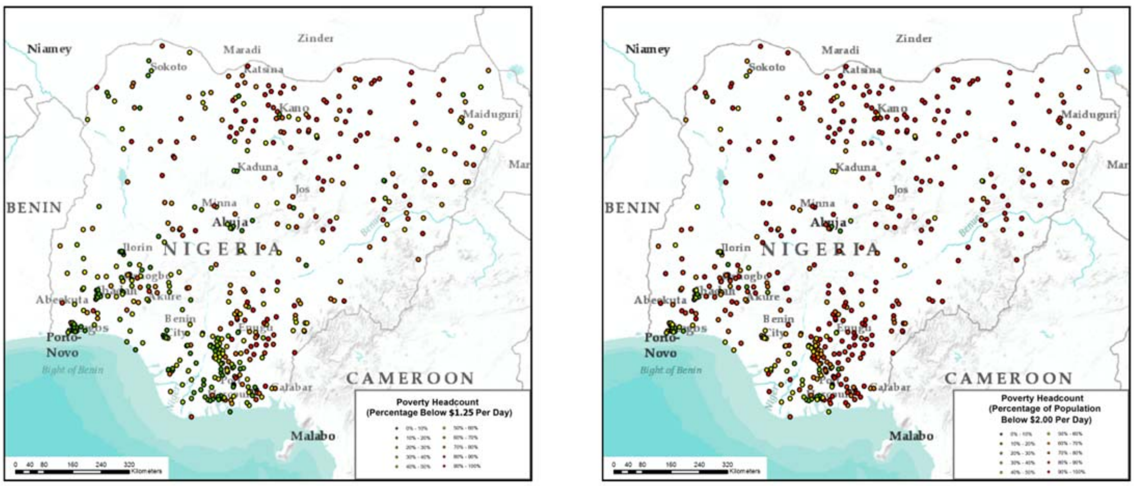

(a) (b)

Figure 4. Input survey clusters for Nigeria displaying the cluster-level consumption-based poverty headcount ratio derived from household surveys, for (a) below $1.25 a day and (b) below $2 a day.

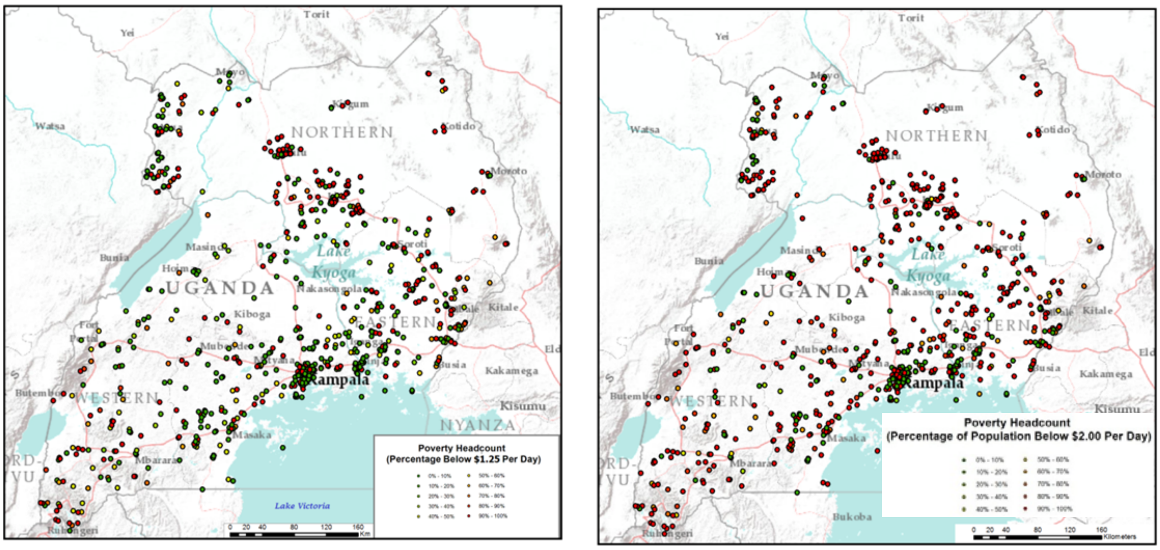

(a) (b)

Figure 5. Input survey clusters for Uganda displaying the cluster-level consumption-based poverty headcount ratio derived from household surveys, for (a) below $1.25 a day and (b) below $2 a day.

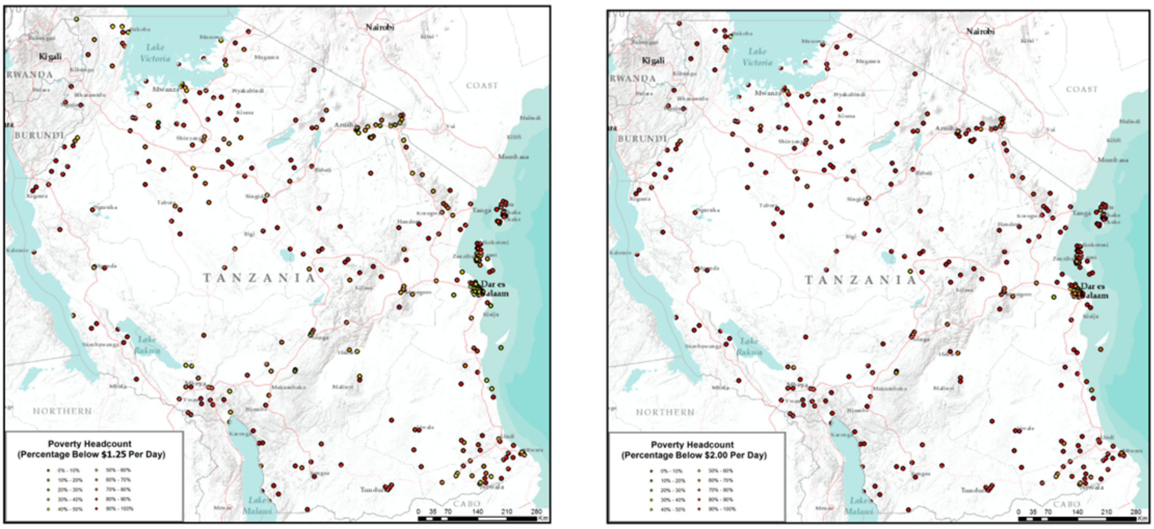

(a) (b)

Figure 6. Input survey clusters for Tanzania displaying the cluster-level consumption-based poverty headcount ratio derived from household surveys, for (a) below $1.25 a day and (b) below $2 a day.

3.0 ASSEMBLING CANDIDATE GEOSPATIAL COVARIATES OF POVERTY

A suite of geospatial covariates were assembled for use in the mapping, focussing on factors likely to have an impact on determining levels of poverty across the six countries considered here. Table 3 provides details on each of the geospatial covariate datasets. The datasets are all provided in differing formats, spatial resolutions, projections and extents. Thus, algorithms were constructed and applied to convert polygon files to gridded datasets and then regrid each gridded dataset to a common 1km spatial resolution grid-frame for use in map production. Figures 7, 8 and 9 show examples of three of the datasets described in table 3 (lights2, evi and access2). While a large suite of data was compiled, not all of the datasets in table 2 were included in the mapping due to differing levels of reliability, relevance and variations in data formats. Many datasets were ultimately left out due to poor spatial resolution and/or categorical inputs that were tested and didn't add sufficient extra information to improve model accuracies.

Category	Dataset Name	Dataset description	Continuous or Categorical	Data Source	Date	Used in mapping
Accessibility	access1	Accessibility to cities with > 50k via all transport methods	Continuous	http://bioval.jrc.ec.europa.eu/products/gam/download.htm	2000	Y
Accessibility	access2	Accessibility via road to towns and cities. Three classes of settlements (size based) and two road types (major and minor) were combined to make a distance weighted layer.	Continuous	Custom product derived from ESRI population datasets and MapAbility road datasets	Population - 2011 Roads – 1980-2012	Y
Population / Urban	aa_pop	Afripop and Asiapop combined dataset	continuous	http://www.afripop.org/ http://www.asiapop.org/	2010	Y
Population / Urban	grump	GRUMP Population Density	continuous	Malaria Atlas Project Master Grids Archive	2010	Y
Population / Urban	gpw	GRUMP population count	continuous	http://sedac.ciesin.columbia.edu/data/set/gpw-v3-population-count/data-download	2000	Y
Crop Suitability	crop61	Map6_61 Suitability of currently available land area for rainfed crops, using maximising crop and technology mix (FGGD)	Quasi Continuous	http://www.fao.org/geonetwork/srv/en/main.home	2007	
Crop Suitability	crop63	Map6_63 Combined suitability of global land area for pasture and rainfed crops (intermediate input level) (FGGD)	Categorical	http://www.fao.org/geonetwork/srv/en/main.home	2007	
Aridity	arid	Mean annual aridity	continuous	http://csi.cgiar.org/Aridity/	1950-2000	Y
Potential Evapotranspiration	pet	Mean annual Potential Evapotranspiration	continuous	http://csi.cgiar.org/Aridity/	1950-2000	Y
Livestock	buffalo	Global buffalo density	Continuous	http://www.fao.org/geonetwork/srv/en/main.home	2005	
Livestock	cattle	Global cattle density	Continuous	http://www.fao.org/geonetwork/srv/en/main.home	2005	
Livestock	goat	Global goat density	Continuous	http://www.fao.org/geonetwork/srv/en/main.home	2005	
Livestock	pig	Global pig density	Continuous	http://www.fao.org/geonetwork/srv/en/main.home	2005	
Livestock	poult	Global poultry density	Continuous	http://www.fao.org/geonetwork/srv/en/main.home	2005	
Livestock	sheep	Global sheep density	Continuous	http://www.fao.org/geonetwork/srv/en/main.home	2005	
Livestock	smrum	Global small ruminant density	Continuous	http://www.fao.org/geonetwork/srv/en/main.home	2005	
Night-time Lights	light1	Global DMSP-OLS Nighttime Lights Time Series	continuous	http://www.ngdc.noaa.gov/dmsp/download.html	2010	
Night-time Lights	lights2	VIIRS Nighttime Lights-2012	continuous	http://www.ngdc.noaa.gov/dmsp/download.html	2012	Y
Ethnicity	ethnic	Geo-referencing of ethnic groups	Categorical	http://www.icr.ethz.ch/data/other/greg	1994	
Elevation	elev	Shuttle Radar Topography Mission (SRTM) Near-global Digital Elevation Models (DEMs)	Continuous	http://webmap.ornl.gov/wcsdown/wcsdown.jsp?dg_id=10008_1	2000	Y
Irrigation	irriga	Global Map of Irrigation Areas version 4.0.1	continuous	http://webmap.ornl.gov/wcsdown/wcsdown.jsp?dg_id=10013_1	2000	
Landcover	lc1	ESA global landcover	Categorical	http://due.esrin.esa.int/globcover/	2009	
Landcover	lc2	IGBP Landcover	Categorical	https://lpdaac.usgs.gov/products/modis_products_table/mcd12q1 ftp://e4ftl01.cr.usgs.gov/MODIS_Composites/MOTA/MCD12Q1.051/2011.01.01/	2011	
Landcover	synmap_1k	Synergetic land cover product (SYNMAP)	categorical	http://webmap.ornl.gov/wcsdown/wcsdown.jsp?dg_id=10024_1	2000	
Protected Areas	protect	United Nations Environment Programme – International Protected areas	Categorical	http://geodata.grid.unep.ch/mod_download/download_geospatial.php?selectedID=1756&newFile=download_pro/WDPA_NATpol2009_po_shp.zip	2009	
MODIS	evi	Mean value for all dates	Continuous	Malaria Atlas Project MasterGrids Archive	2001-2005	Y
MODIS	lst	Mean value for all dates	Continuous	Malaria Atlas Project MasterGrids Archive	2001-2005	Y
MODIS	midir	Mean value for all dates	Continuous	Malaria Atlas Project MasterGrids Archive	2001-2005	Y

Water Bodies	water	WWF Global Lakes and Wetlands Database	Categorica	http://worldwildlife.org/pages/global-lakes-and-wetlands-database	2012	

Table 3. Details of geospatial covariates assembled for use in mapping.

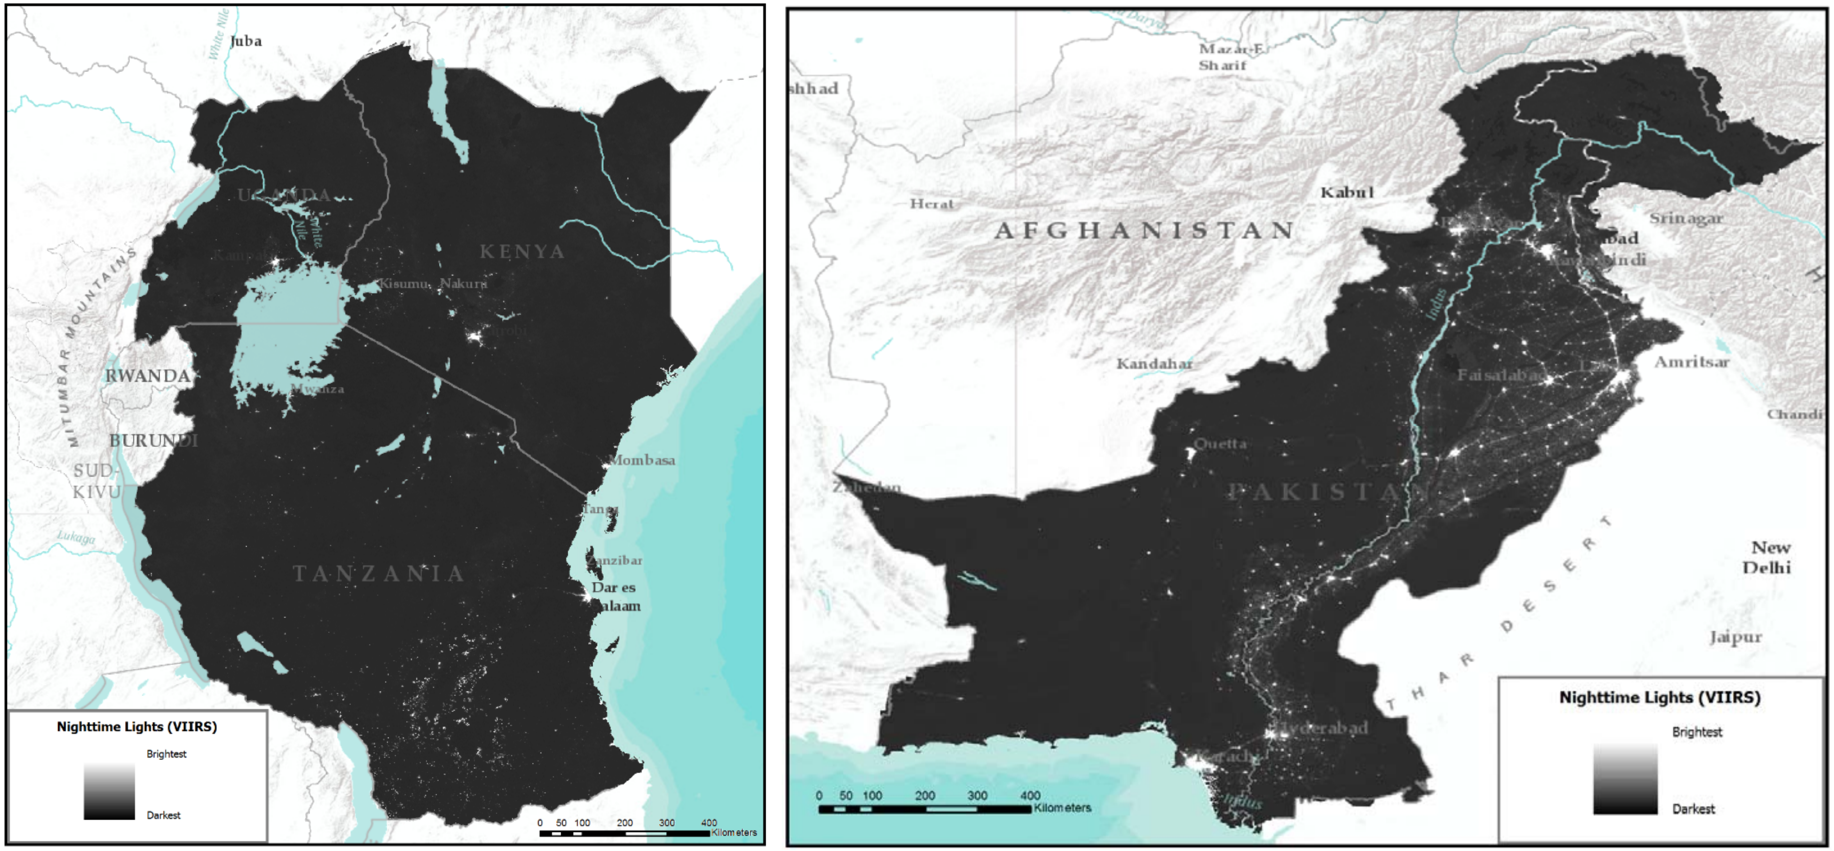

(a) **(b)**

Figure 7. Examples of geospatial covariate datasets: Annual composite nighttime light satellite imagery from the VIIRS satellite sensor ('lights2') for (a) East Africa and (b) Pakistan.

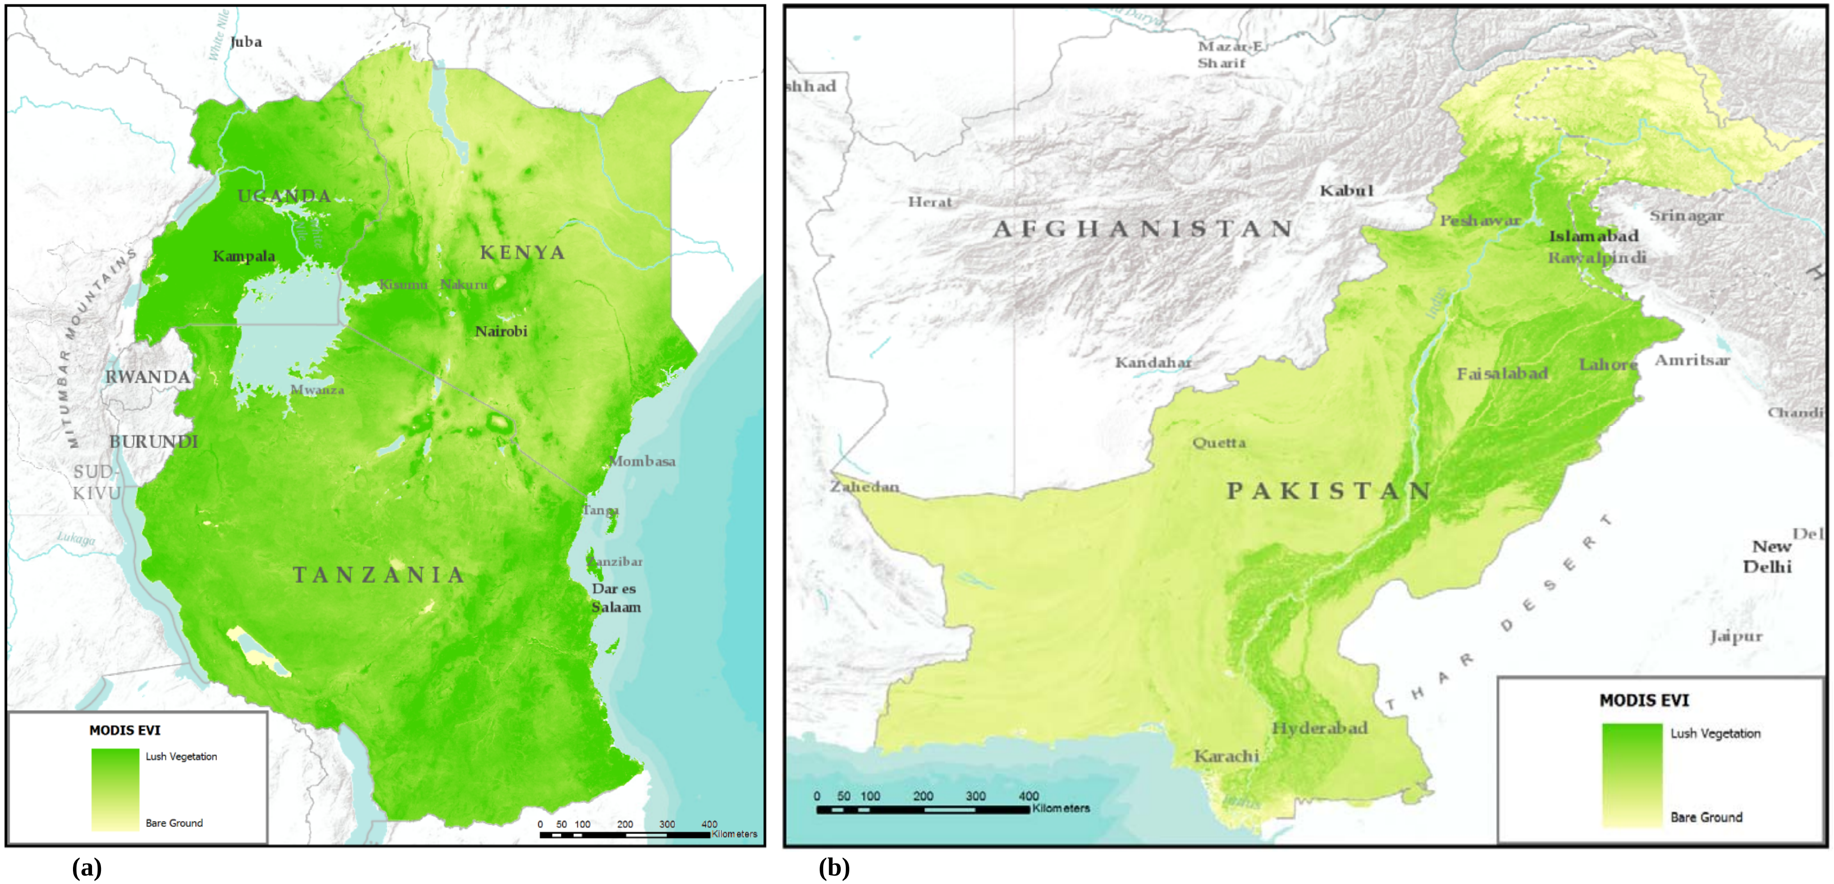

Figure 8. Examples of geospatial covariate datasets: Mean annual enhanced vegetation index ('evi') from the MODIS satellite sensor for (a) East Africa and (b) Pakistan.

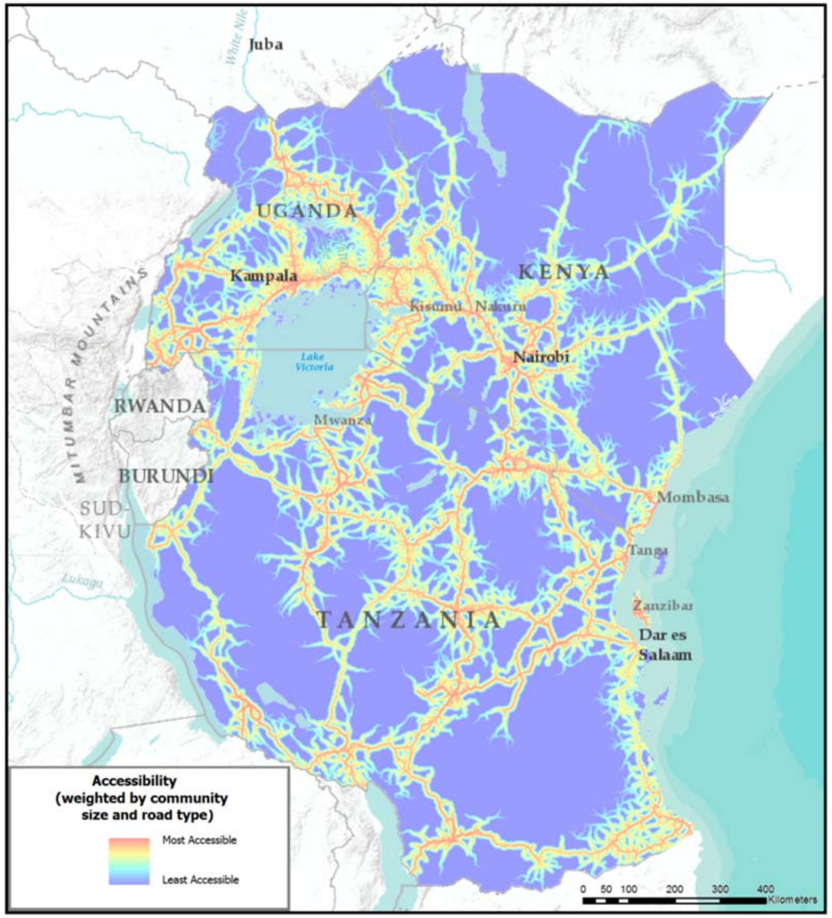

(a)

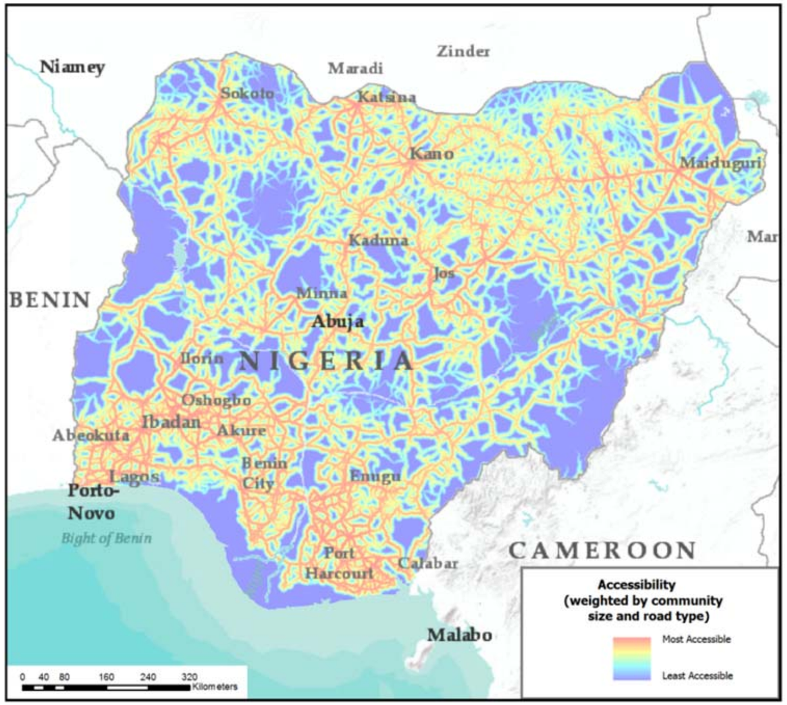

(b)

Figure 9. Examples of geospatial covariate datasets: Relative accessibility to major population centers for (a) East Africa and (b) Nigeria.

4.0 CREATION OF PILOT POVERTY MAPS

Using either the MPI headcount or the <$1.25 or <$2 a day headcount as our test variable, we have implemented a model-based geostatistical framework to generate pilot poverty maps at 1x1 km resolution for Kenya, Uganda, Tanzania, Nigeria, Malawi and Pakistan. Here we describe the exploratory analysis, model formulation and validation, and present the output mapped surfaces.

4.1 METHODS

4.1.1 Model structure

Our initial model structure is a class of generalized linear mixed model, with an approximation of a multivariate Normal random field (i.e. a Gaussian Process) used as a spatially autocorrelated random effect term. This family of models derives from a body of theory knows as model-based geostatistics [23,24]. Below we present the methods used for modeling MPI headcount ratio, but the methods for the consumption-based modeling are the same.

The MPI headcount ratio (proportion of individuals considered 'poor' according to the multi-dimensional poverty index) $MPIh(x_i)$ (replace this with $CONSh(x_i)$ for consumption-based metrics), at each location in the country of interest x_i was modeled as a transformation $g(\cdot)$ of a spatially structured field superimposed with additional random variation $\epsilon(x_i)$. The count of individuals considered poor N_i^+ from the total sample of N_i in each survey cluster was modeled as a conditionally independent binomial variate given the unobserved underlying $MPIh(x_i)$ value. The spatial component was represented by a stationary Gaussian process $f(x_i, t_i)$ with mean μ and covariance C. The unstructured component $\epsilon(x_i)$ was represented as Gaussian with zero mean and variance V. Both the inference and prediction stages were coded using the INLA framework [25], primarily in R [26].

4.1.1.1 Mean and covariance definition

The mean component μ was modelled as a linear function of *n*=12 environmental covariates, $\mu = \boldsymbol{\beta}\mathbf{X}$,where $\mathbf{X} = 1, X_1(x), ..., X_n(x)$ was a vector consisting of a constant and the covariates indexed by spatial location x, and $\boldsymbol{\beta} = \beta_0, \beta_1, ..., \beta_n$ was a corresponding vector of regression coefficients. Each covariate was converted to *z*-scores before analysis. In this pilot stage, we took an inclusive approach to covariate selection, with most of the assembled variables being included in the model. As our library of covariates becomes more refined, we will implement more formal model selection procedures will be implemented in future work to identify optimal covariate suites for inclusion.

Covariance between spatial locations was modeled using a Matern covariance function C:

$$C\left(d(x_i;x_j)\right)=\sigma^2\frac{1}{\Gamma(\upsilon)2^{\upsilon-1}}\left(\sqrt{2\upsilon}\frac{d(x_i;x_j)}{\rho}\right)^{\upsilon}K_{\upsilon}\left(\sqrt{2\upsilon}\frac{d(x_i;x_j)}{\rho}\right)$$

Where $d(x_i;x_j)$ is the geographical separation between two points; σ,υ,ρ are parameters of the covariance function defining, respectively, its amplitude, degree of differentiability, and scale; K_{υ} is the modified Bessel function of the second kind of order υ, and Γ is the gamma function [27,28].

4.1.2 Model implementation and output

Bayesian inference was implemented using the INLA algorithm to generate approximations of the marginal posterior distributions of the outcome variable $MPIh(x_i)$ or $CONSh(x_i)$ at each location on a regular 1 × 1 km spatial grid across the country of interest and of the unobserved parameters of the mean, covariance function and Gaussian random noise component. At each location, the posterior distribution was summarized using the posterior mean as a point estimate, and the posterior inter-quartile range as a measure of model precision. Maps were generated of each of these metrics in ArcGIS 10.2.

4.1.3 Validation

The predictive performance of the model was assessed via out-of-sample validation. We implemented a ten-fold hold-out procedure whereby 10% of the data points were randomly withdrawn from the dataset, the model run in full using the remaining 90% of data, and the predicted values at the locations of the hold-out data compared to their observed values. This was repeated ten times without replacement such that every data point was held out once across the ten validation runs. Standard validation statistics were computed as measures of model precision (root mean square error, mean absolute error), bias (mean error), and correlation between observed and predicted. We also generated a scatter plot of observed versus predicted values for visualization purposes.

5.0 RESULTS

A sample of results from both the MPI and consumption-based metric mapping are presented here to highlight the key results of analyses undertaken. More detailed results can be obtained through contacting the authors, and mapped outputs can be freely downloaded through the WorldPop project website (www.worldpop.org.uk).

5.1 MPI RESULTS

5.1.1 Model fit and validation statistics: East Africa MPI

	mean	sd	q0.025	q0.5	q0.975
Fixed effects					
Intercept	1.113	0.709	-0.278	1.113	2.503
aa_pop	**-0.008**	**0.003**	**-0.014**	**-0.008**	**-0.002**
access1	**0.004**	**0.000**	**0.004**	**0.004**	**0.005**
access2	0.031	0.020	-0.008	0.031	0.071
arid	**-0.156**	**0.078**	**-0.308**	**-0.156**	**-0.003**
elev	0.033	0.092	-0.147	0.033	0.214
evi	-0.206	0.179	-0.556	-0.206	0.145
gpw	0.003	0.003	-0.004	0.003	0.010
grump	0.004	0.003	-0.002	0.004	0.009
lights2	**-0.106**	**0.012**	**-0.129**	**-0.106**	**-0.083**
lst	0.084	0.380	-0.660	0.084	0.829
midir	-0.093	0.073	-0.237	-0.093	0.050
pet	0.261	0.320	-0.366	0.261	0.888
Random effects					
Theta1 (log)	**-4.067**	**0.037**	**-4.139**	**-4.067**	**-3.992**
Theta2 (log)	**2.604**	**0.039**	**2.525**	**2.605**	**2.678**
Precision	**1.146**	**0.034**	**1.084**	**1.145**	**1.216**

Table 4. Summary of posterior distributions of regression coefficients for 12 covariates and three parameters of the covariance function for the East Africa MPI model. sd = posterior standard deviation; q0.025, q0.5, q0.975 are the posterior 2.5th, 50th (median) and 97.5th percentiles. Fixed effects are listed in descending order of the magnitude of the posterior mean effect. Coefficients whose 95% credible interval do not include zero are highlighted in bold.

Table 4 summarizes the East Africa MPI model fit: each parameter is listed and summaries of its posterior distribution are given. The geospatial covariate names referred to are documented and described in section 3.0. Four of the 12 coefficients were ‘significant’ in that their 95% credible interval did not span zero. Of the ‘non-significant’ coefficients, however, several were large in magnitude. The largest effects were associated with some of the environmental variables, such as the enhanced vegetation index (evi), aridity (arid) and potential evapotranspiration (pet). The satellite nighttime lights also showed a relatively large effect.

Validation metric	Value
Deviance Information Criteria	7965.27
Mean Square Error	0.037
Root Mean Square Error	0.193
Mean Error	-0.003
Absolute Error	0.15
Hold out correlation	0.972

Table 5. Validation statistics for East Africa

Table 5 displays validation statistics derived from the ten-fold out-of-sample procedure. The model is essentially unbiased (mean error = -0.003) indicating no overall tendency to over- or under-predict the MPI headcount ratio. Mean absolute error was 0.15, which should be interpreted directly in the same units as the outcome variable (i.e. MPI is a proportion between zero and one). The correlation between predicted and actual values was 0.972, indicating an excellent degree of linear association. This can also be seen in the scatterplot in figure 10.

5.1.2 Mapped surfaces: East Africa MPI

Our pilot predicted map of the MPI headcount ratio for East Africa (Kenya, Uganda and Tanzania) is displayed in Figure 11a. At the macro-scale, we predict regions with the lowest rates of multi-dimensional poverty associated with Nairobi, and the extended central (and most densely populated) areas of Kenya surrounding the capital and extending westward throughout the highlands and towards Lake Victoria. The southern coastal region from Mombasa northwards to Malindi is also relatively less poor, as are the coastal regions of Tanga and Dar es Salaam and the island of Unguja in Zanzibar. In Uganda, the Kampala region represents the area with the predicted lowest proportions of the population living in poverty. In contrast, the more arid areas of northern and eastern Kenya, the northern areas of Uganda and many parts of rural Tanzania are associated with the highest rates of multi-dimensional poverty. Figure 11b provides one indication of the geographically varying precision of our model output: the 95% credible interval for each pixel. Small values are associated with more precise, less uncertain, model predictions. This precision is affected by (i) the local density of survey points (with more points leading to higher precision) but also (ii) the local heterogeneity of observed values. This means that, for example, areas in the arid northwest of Kenya, whilst relatively sparsely surveyed, are predicted with high precision because observed values have low variance (consistently reporting very high rates of poverty).

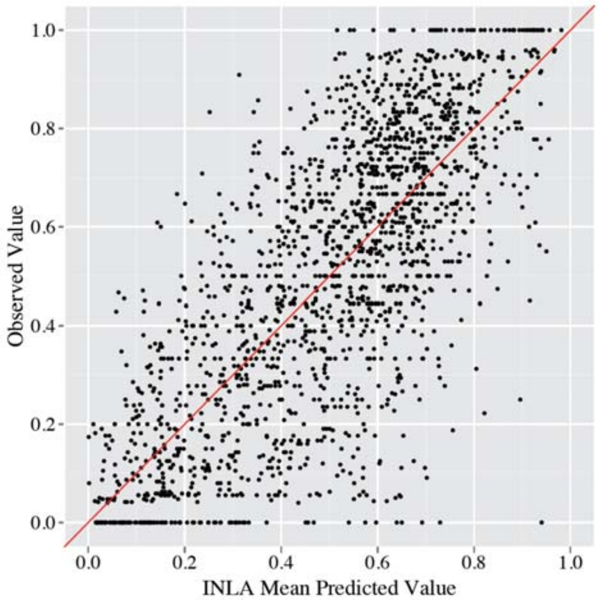

Figure 10. Scatterplot of observed versus predicted values for East Africa. The one-to-one line is shown in red.

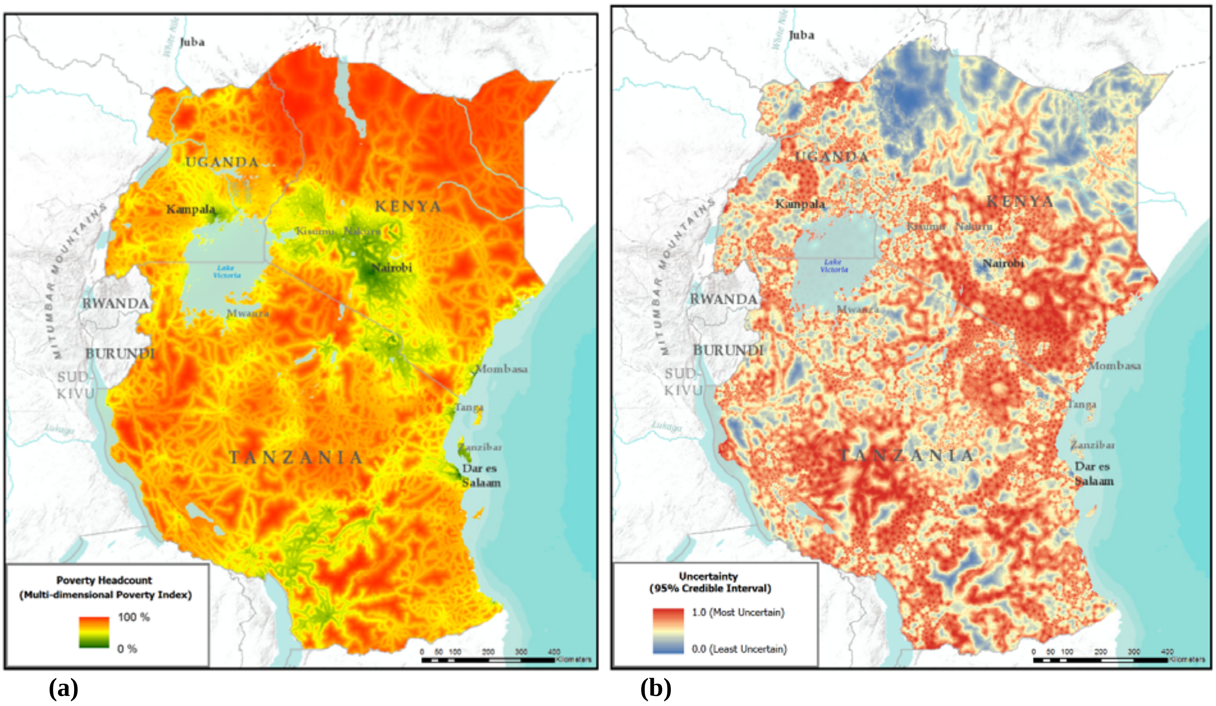

Figure 11. (a) Predicted map of the MPI headcount ratio for Kenya, Uganda and Tanzania. This displays the mean value of the predictive posterior distribution at each 1x1km pixel. Major waterbodies and city names are also overlaid for context. (b) The precision of the model output for East Africa as measured using the 95% credible interval. Major waterbodies and city names are also overlaid for context.

5.1.3 Model fit and validation statistics: Pakistan MPI

	mean	sd	q0.025	q0.5	q0.975
Fixed effects					
Intercept	**2.561**	**0.484**	**1.613**	**2.560**	**3.512**
aa_pop	-0.002	0.004	-0.010	-0.002	0.006
access1	**6.050**	**1.472**	**3.164**	**6.050**	**8.938**
access2	**2.031**	**0.604**	**0.847**	**2.031**	**3.216**
arid	-0.336	0.413	-1.146	-0.336	0.475
elev	0.066	0.549	-1.008	0.065	1.146
evi	**2.467**	**0.506**	**1.475**	**2.467**	**3.461**
gpw	0.001	0.006	-0.011	0.001	0.014
grump	0.000	0.004	-0.009	0.000	0.008
lights2	**-0.080**	**0.017**	**-0.113**	**-0.080**	**-0.047**
lst	**-2.289**	**1.013**	**-4.277**	**-2.289**	**-0.301**
midir	**0.608**	**0.144**	**0.325**	**0.608**	**0.891**
pet	-0.429	0.761	-1.925	-0.427	1.062
Random effects					
Theta1 (log)	**-5.778**	**0.055**	**-5.886**	**-5.778**	**-5.669**
Theta2 (log)	**4.588**	**0.070**	**4.462**	**4.583**	**4.736**
Precision	**3.556**	**0.231**	**3.145**	**3.540**	**4.052**

Table 6. Summary of posterior distributions of regression coefficients for 12 covariates and three parameters of the covariance function for the Pakistan MPI model. sd = posterior standard deviation; q0.025, q0.5, q0.975 are the posterior 2.5th, 50th (median) and 97.5th percentiles. Fixed effects are listed in descending order of the magnitude of the posterior mean effect. Coefficients whose 95% credible interval do not include zero are highlighted in bold.

Table 6 summarizes the model fit for Pakistan: each parameter is listed and summaries of its posterior distribution are given. The geospatial covariate names referred to are documented and described in section 3.0. Seven of the 12 coefficients were 'significant' in that their 95% credible interval did not span zero. Of the 'non-significant' coefficients, however, several were large in magnitude. The largest effects were associated with geographical isolation (access1), and environmental measures, including enhanced vegetation index (evi) and land surface temperature (lst). Each of these covariates likely aided in distinguishing the poorer mountainous areas from the rest of the country.

Validation metric	Value
Deviance Information Criteria	3386.45
Mean Square Error	0.048
Root Mean Square Error	0.218
Mean Error	-0.003
Absolute Error	0.175
Hold out correlation	0.921

Table 7. Validation statistics for Pakistan

Table 7 displays validation statistics derived from the ten-fold out-of-sample procedure. The model is essentially unbiased (mean error = -0.003) indicating no overall tendency to over- or under-predict the MPI headcount ratio. Mean absolute error was 0.175, which should be interpreted directly in the same units as the outcome variable (i.e. MPI is a proportion between zero and one). The correlation between predicted and actual values was 0.921, indicating an excellent degree of linear association. This can also be seen in the scatterplot in Figure 12.

5.1.4 Mapped surfaces: Pakistan MPI

Our pilot predicted map of the MPI headcount ratio for Pakistan is displayed in Figure 13a. At the macro-scale, we predict regions with the lowest rates of multi-dimensional poverty associated with the northern regions in and around Lahore and Islamabad, the cities of Hyderabad and Karachi in the south and the southern border region with Iran. In contrast, the remote mountainous and central/south rural regions have the highest predicted proportions of their populations classed as in poverty. Figure 13b provides one indication of the geographically varying precision of our model output: the 95% credible interval for each pixel. Small values are associated with more precise, less uncertain, model predictions. As described above, this precision is affected by (i) the local density of survey points (with more points leading to higher precision) but also (ii) the local heterogeneity of observed values.

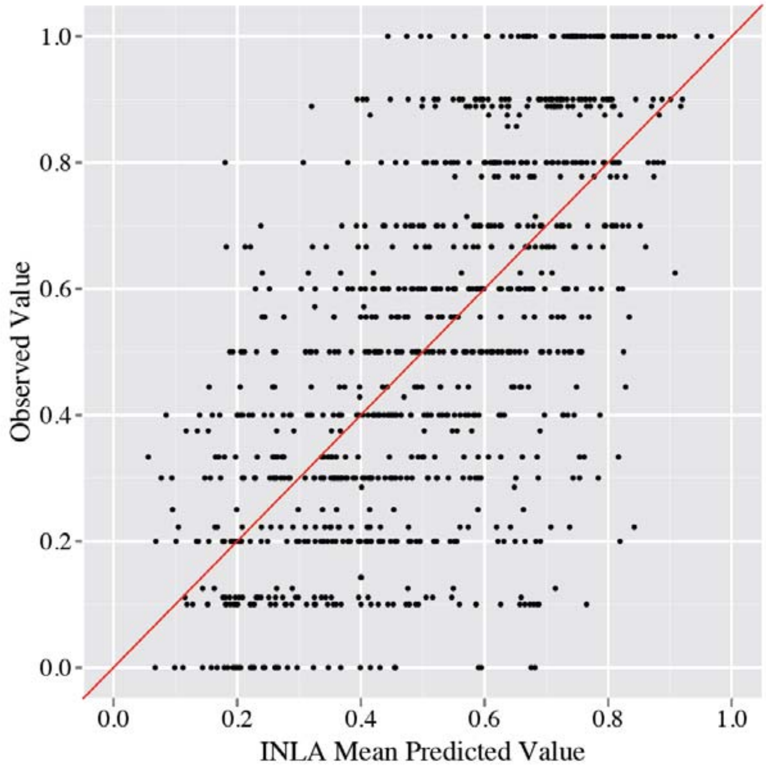

Figure 12. Scatterplot of observed versus predicted values for Pakistan. The one-to-one line is shown in red.

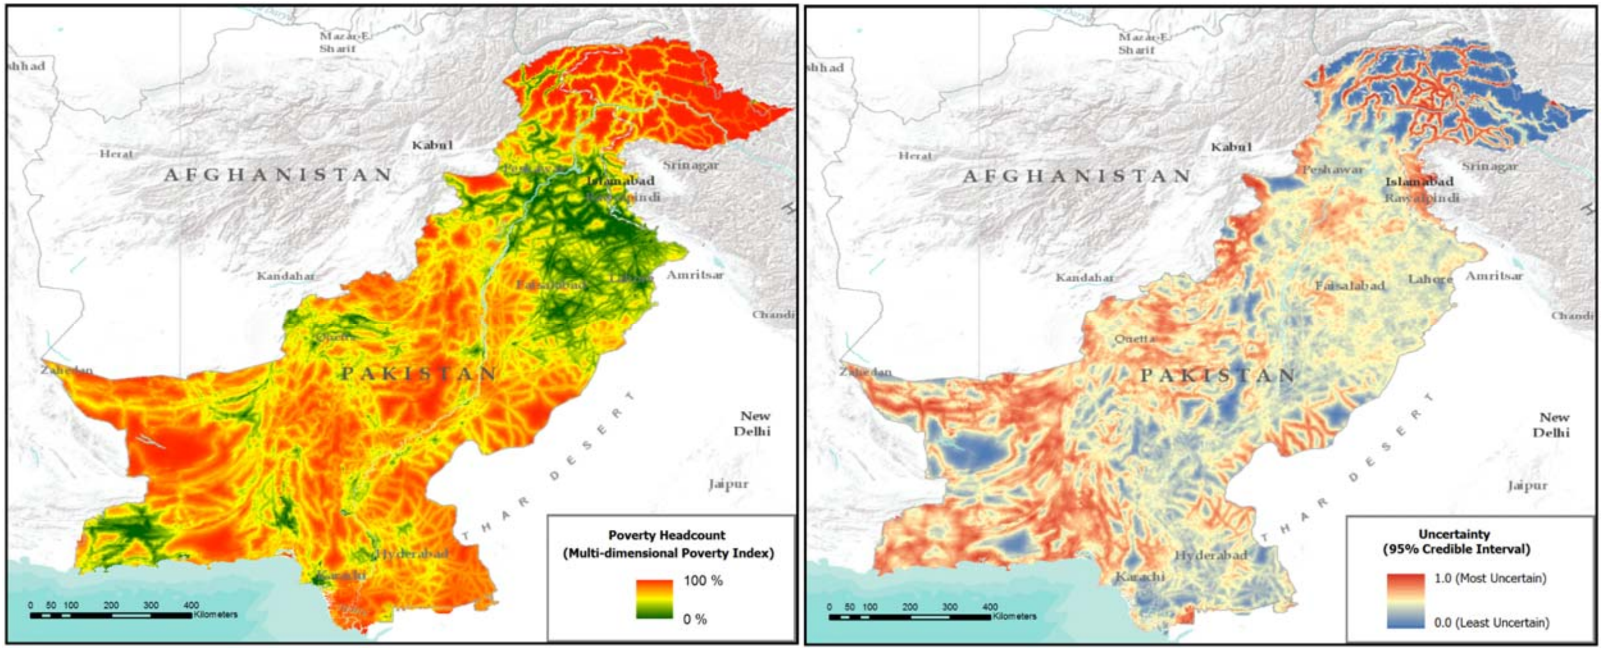

(a) **(b)**

Figure 13. (a) Predicted map of the MPI headcount ratio for Pakistan. This displays the mean value of the predictive posterior distribution at each 1x1km pixel. Major waterbodies and city names are also overlaid for context; (b) The precision of the model output for Pakistan as measured using the 95% credible interval. Major waterbodies and city names are also overlaid for context.

5.2 CONSUMPTION-BASED MAPPING RESULTS

Results for the consumption-based mapping follow a similar format to those presented above for MPI in section 5.1. Overall, model fits were weaker for the consumption-based mapping. This is unsurprising, given that the MPI contains more environmentally-sensitive components than consumption. However, output statistics for the consumption-based model fits do still indicate a similar level of mapping accuracy to those achieved by existing census-based approaches used by the World Bank, and formal comparisons of the two approaches are underway (see section 7). Moreover, for some countries analysed here, it is clear that the vast majority of the rural populations live below the $2 a day line, and therefore, little spatial variation exists to for capturing with model fitting (e.g. see figure 6b for input data for Tanzania).

5.2.1 Model fit and validation statistics: Consumption-based mapping

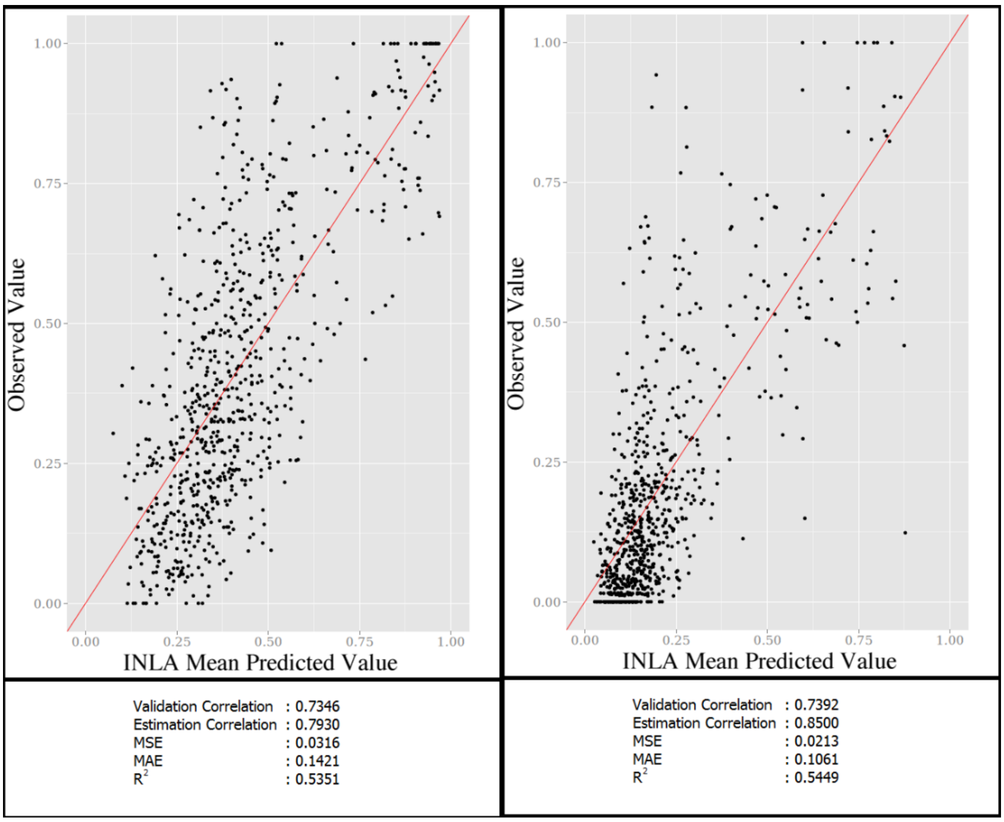

(a) (b)

Figure 14. Scatterplots (the one-to-one line is shown in red) and validation statistics for Malawi, for (a) below $1.25 a day, and (b) below $2 a day.

The models presented in figure 14 for Malawi are essentially unbiased (mean square error = -0.02, 0.03) indicating no overall tendency to over- or under-predict the poverty headcount ratio.

Mean absolute errors were between 0.11 and 0.14, which should be interpreted directly in the same units as the outcome variable (i.e. poverty headcount is a proportion between zero and one). The correlation between predicted and actual values was around 0.73, indicating a good degree of linear association. This can also be seen in the scatterplot in figure 14.

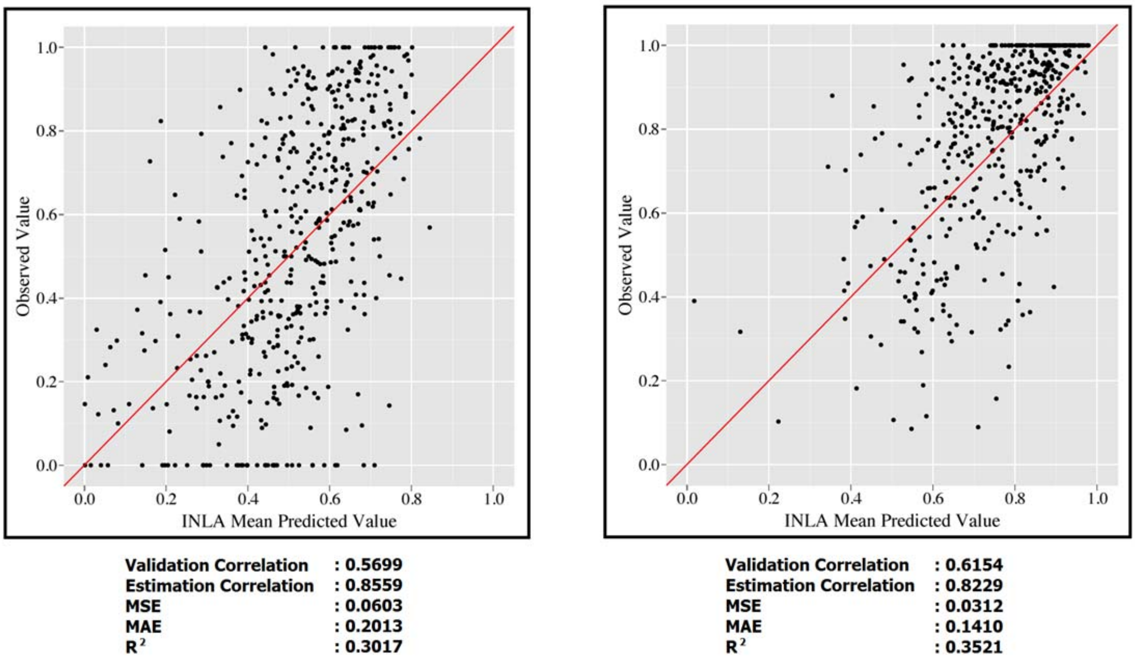

(a) **(b)**

Figure 15. Scatterplots (the one-to-one line is shown in red) and validation statistics for Nigeria, for (a) below \$1.25 a day, and (b) below \$2 a day.

The models presented in figure 15 for Nigeria are essentially unbiased (mean square error = -0.06, 0.03) indicating no overall tendency to over- or under-predict the poverty headcount ratio. Mean absolute errors were between 0.14 and 0.2, which should be interpreted directly in the same units as the outcome variable (i.e. poverty headcount is a proportion between zero and one). The correlations between predicted and actual values were between 0.57 and 0.62, indicating a generally good degree of linear association. This can also be seen in the scatterplot in figure 15.

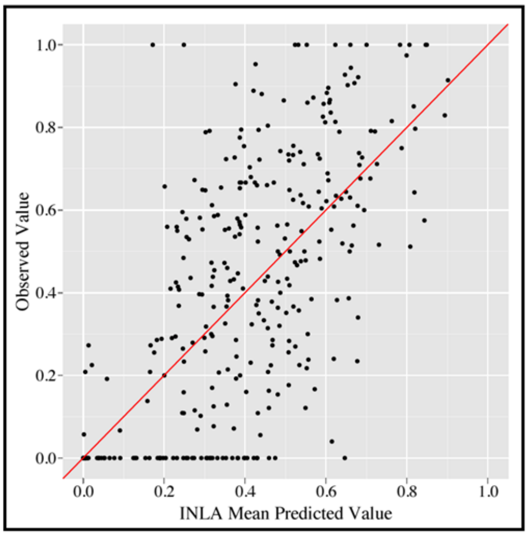

Validation Correlation	**: 0.4809**
Estimation Correlation	**: 0.7298**
MSE	**: 0.0559**
MAE	**: 0.1883**
R^2	**: 0.2370**

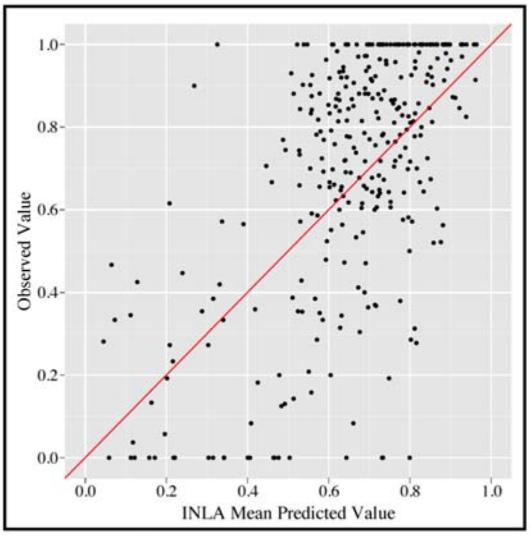

Validation Correlation	**: 0.6146**
Estimation Correlation	**: 0.7007**
MSE	**: 0.0573**
MAE	**: 0.1918**
R^2	**: 0.3713**

(a) **(b)**

Figure 16. Scatterplots (the one-to-one line is shown in red) and validation statistics for Uganda, for (a) below \$1.25 a day, and (b) below \$2 a day.

The models presented in figure 16 for Uganda are essentially unbiased (mean square error = -0.06) indicating no overall large tendency to over- or under-predict the poverty headcount ratio. Mean absolute errors were between 0.18 and 0.19, which should be interpreted directly in the same units as the outcome variable (i.e. poverty headcount is a proportion between zero and one). The correlations between predicted and actual values were between 0.48 and 0.61, indicating a generally good degree of linear association. This can also be seen in the scatterplot in figure 16.

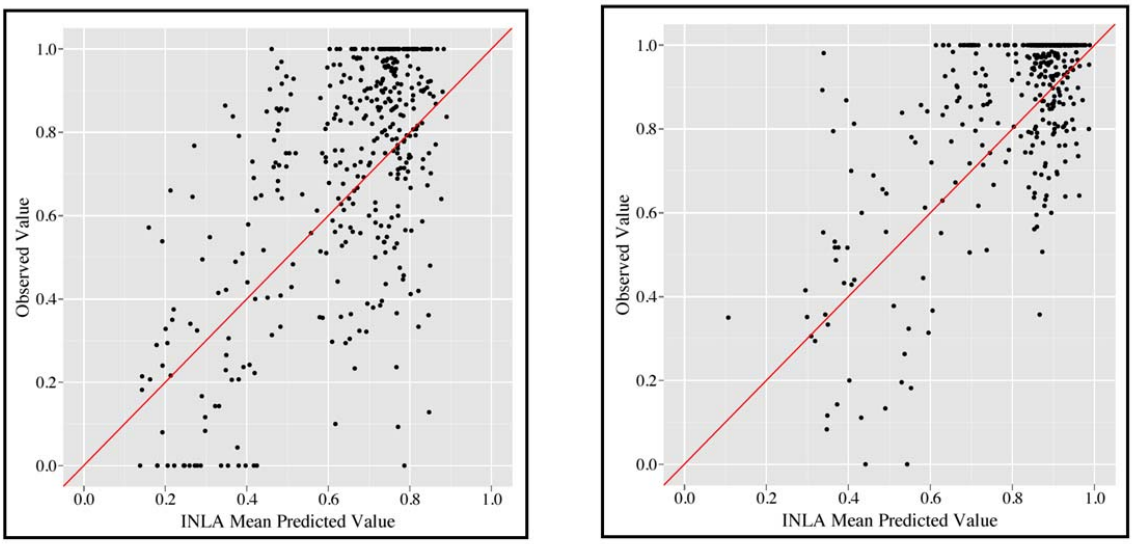

Validation Correlation	**: 0.4878**
Estimation Correlation	**: 0.8617**
MSE	**: 0.0528**
MAE	**: 0.1908**
R^2	**: 0.2418**

Validation Correlation	**: 0.4816**
Estimation Correlation	**: 0.8620**
MSE	**: 0.0266**
MAE	**: 0.1287**
R^2	**: 0.2251**

(a) **(b)**

Figure 17. Scatterplots (the one-to-one line is shown in red) and validation statistics for Tanzania, for (a) below $1.25 a day, and (b) below $2 a day.

The models presented in figure 17 for Tanzania are essentially unbiased (mean square error = -0.05, 0.03) indicating no overall large tendency to over- or under-predict the poverty headcount ratio. Mean absolute errors were between 0.13 and 0.19, which should be interpreted directly in the same units as the outcome variable (i.e. poverty headcount is a proportion between zero and one). The correlations between predicted and actual values were around 0.48, indicating a satisfactory degree of linear association. This can also be seen in the scatterplot in figure 17. The high proportion of surveys showing 100% of households below the $1.25 and $2 a day lines made model fitting challenging though.

5.2.2 Mapped surfaces: Consumption-based mapping

Figures 18, 19, 20 and 21 show the output <$1.25 and <$2 a day poverty maps and associated 95% credible interval maps for Malawi, Nigeria, Uganda and Tanzania. As with the MPI maps, there is again a common factor of higher uncertainty in (i) urban areas where input survey values vary substantially over small distances, and also (ii) in some rural areas where cluster data are lacking to inform the model. Moreover, urban populations are again shown to be consistently less poor on average than their rural counterparts. Finally, for Nigeria and Tanzania in particular, the vast majority of rural populations surveyed are below the $2 a day line, making the outputs maps for this poverty line relatively homogenous.

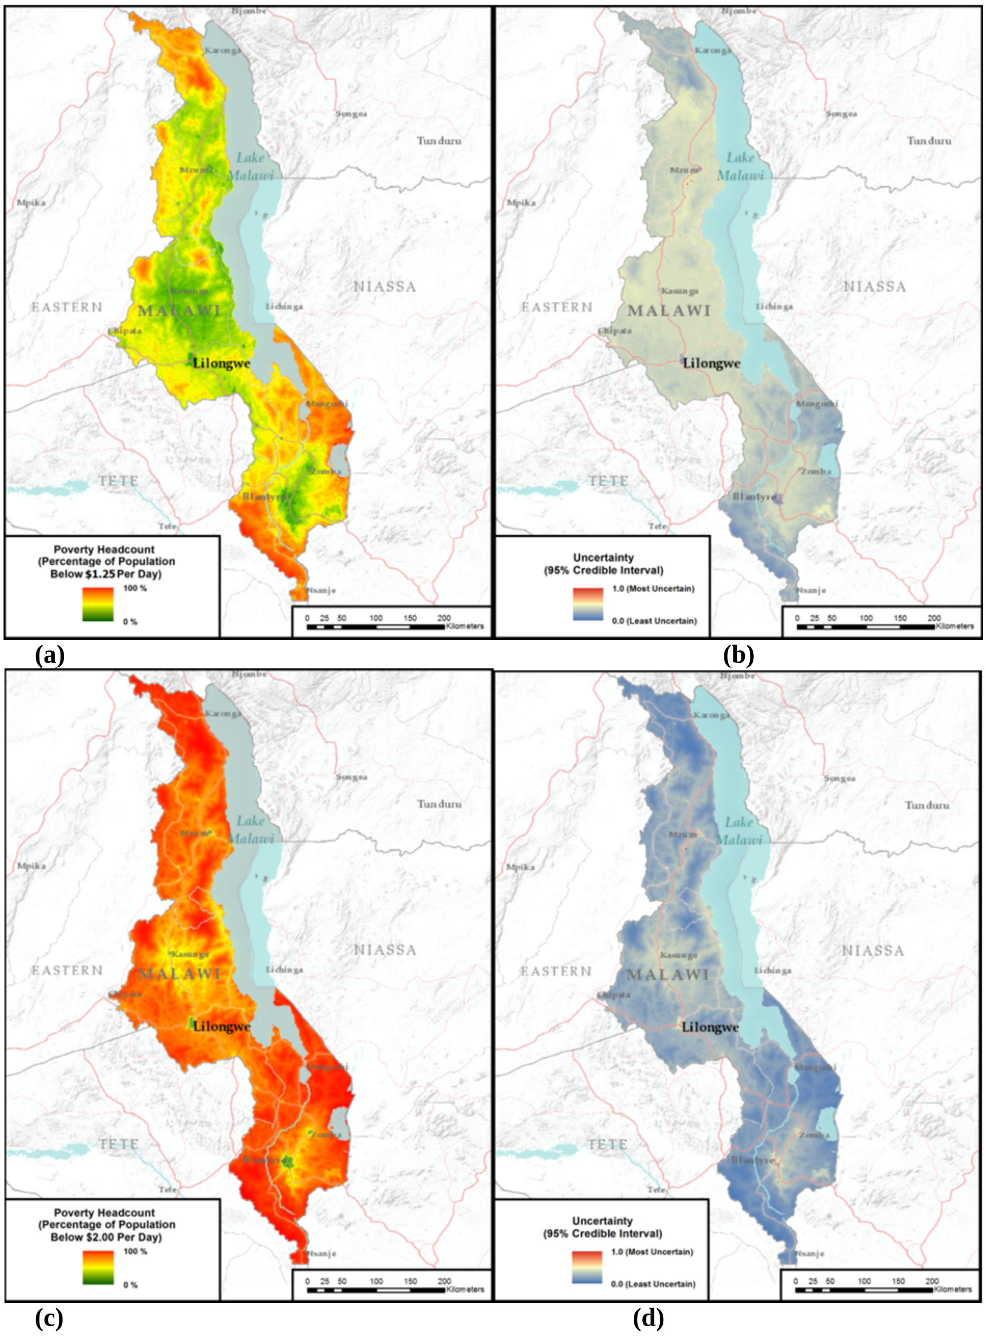

Figure 18. (a) Predicted map of the below $1.25 a day headcount ratio for Malawi. This displays the mean value of the predictive posterior distribution at each 1x1km pixel. Major waterbodies and city names are also overlaid for context. (b) The precision of the <$1.25 a day model output for Malawi as measured using the 95% credible interval. Major waterbodies and city names are also overlaid for context. (c) Predicted map of the below $2 a day headcount ratio for Malawi. This displays the mean value of the predictive posterior distribution at each 1x1km pixel. Major waterbodies and city names are also overlaid for context. (d) The precision of the <$2 a day model output for Malawi as measured using the 95% credible interval. Major waterbodies and city names are also overlaid for context.

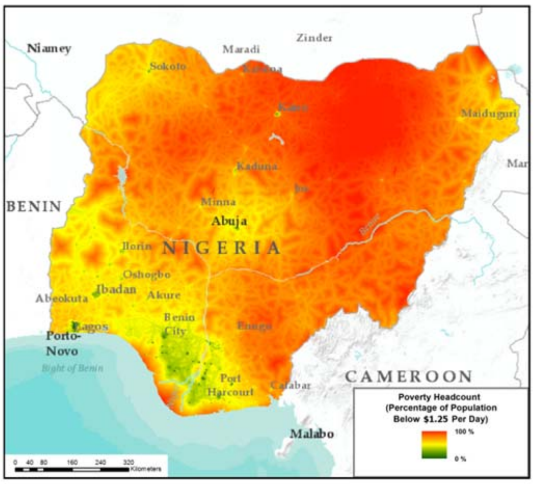

(a)

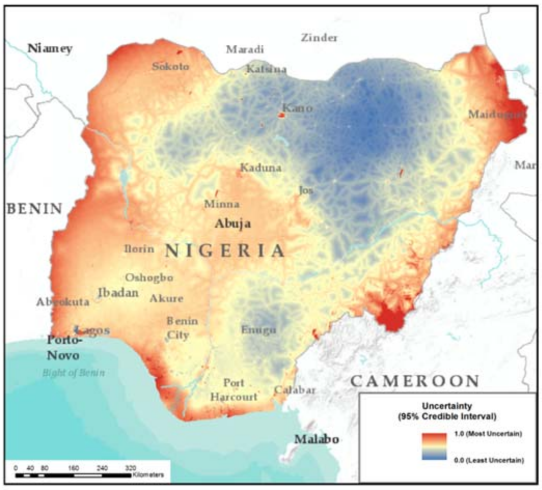

(b)

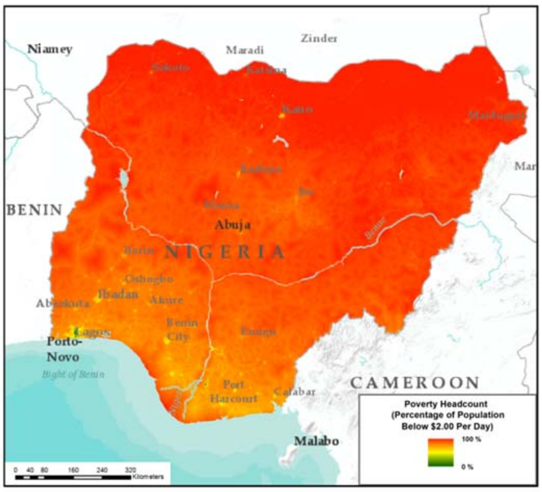

(c)

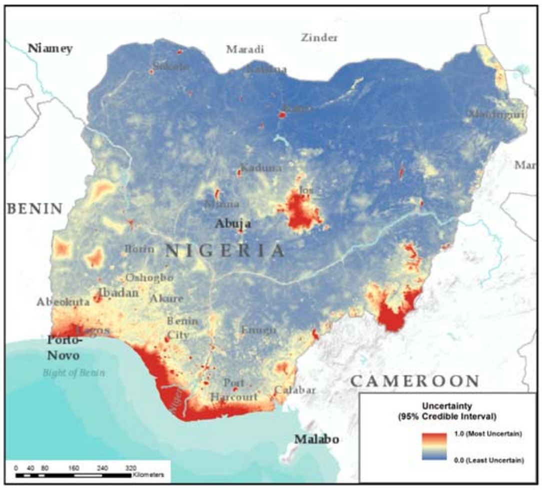

(d)

Figure 19. (a) Predicted map of the below $1.25 a day headcount ratio for Nigeria. This displays the mean value of the predictive posterior distribution at each 1x1km pixel. Major waterbodies and city names are also overlaid for context. (b) The precision of the <$1.25 a day model output for Nigeria as measured using the 95% credible interval. Major waterbodies and city names are also overlaid for context. (c) Predicted map of the below $2 a day headcount ratio for Nigeria. This displays the mean value of the predictive posterior distribution at each 1x1km pixel. Major waterbodies and city names are also overlaid for context. (d) The precision of the <$2 a day model output for Nigeria as measured using the 95% credible interval. Major waterbodies and city names are also overlaid for context.

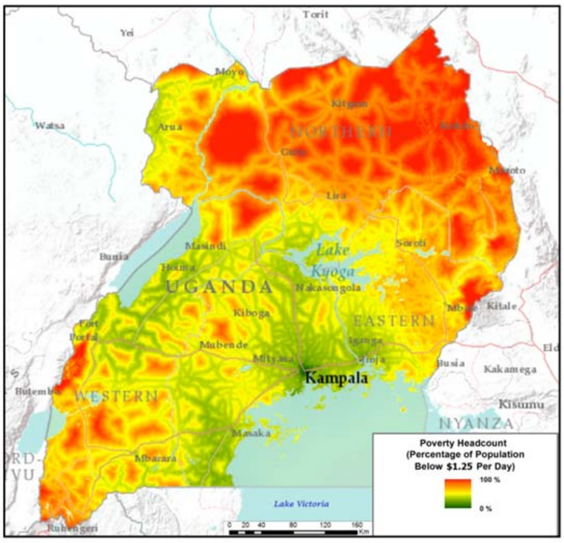

(a)

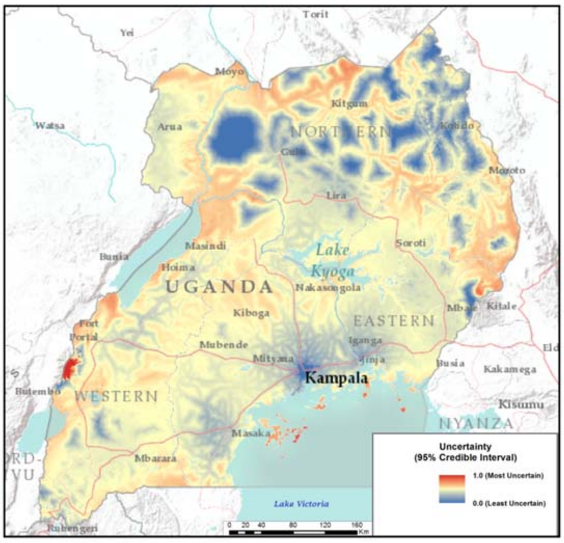

(b)

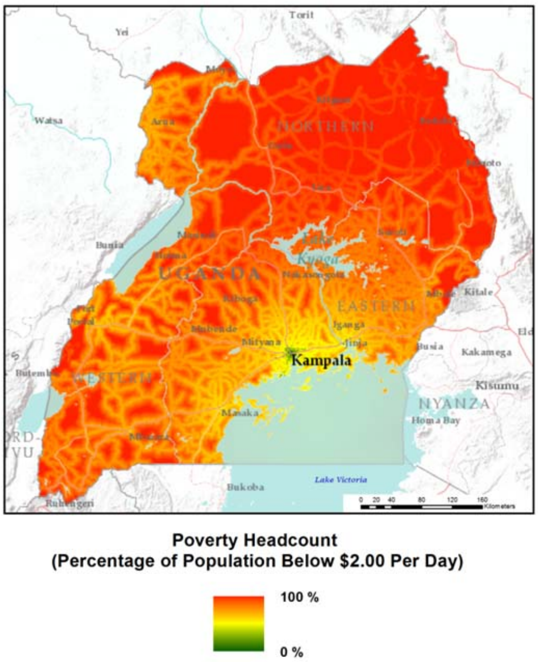

(c)

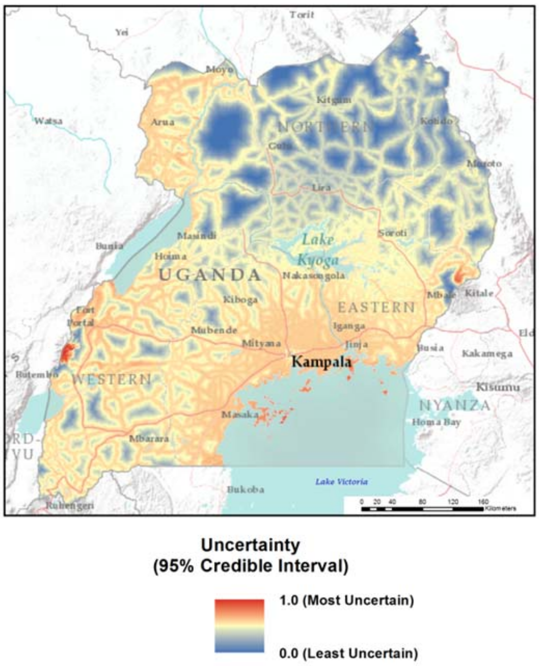

(d)

Figure 20. (a) Predicted map of the below $1.25 a day headcount ratio for Uganda This displays the mean value of the predictive posterior distribution at each 1x1km pixel. Major waterbodies and city names are also overlaid for context. (b) The precision of the <$1.25 a day model output for Uganda as measured using the 95% credible interval. Major waterbodies and city names are also overlaid for context. (c) Predicted map of the below $2 a day headcount ratio for Uganda. This displays the mean value of the predictive posterior distribution at each 1x1km pixel. Major waterbodies and city names are also overlaid for context. (d) The

precision of the <$2 a day model output for Uganda as measured using the 95% credible interval. Major waterbodies and city names are also overlaid for context.

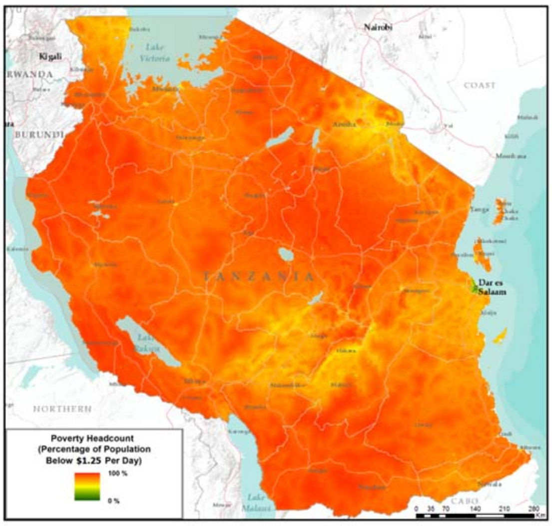

(a)

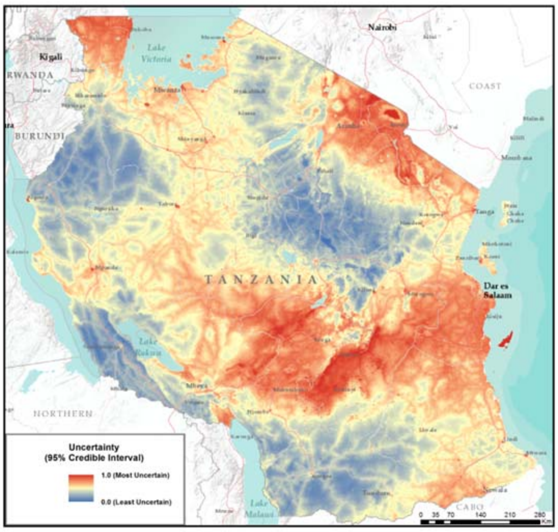

(b)

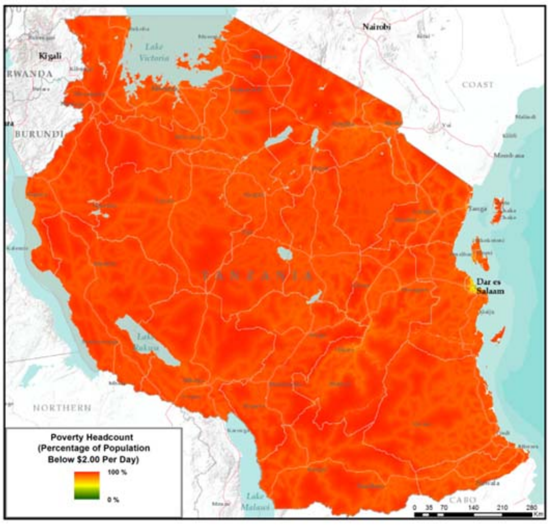

(c)

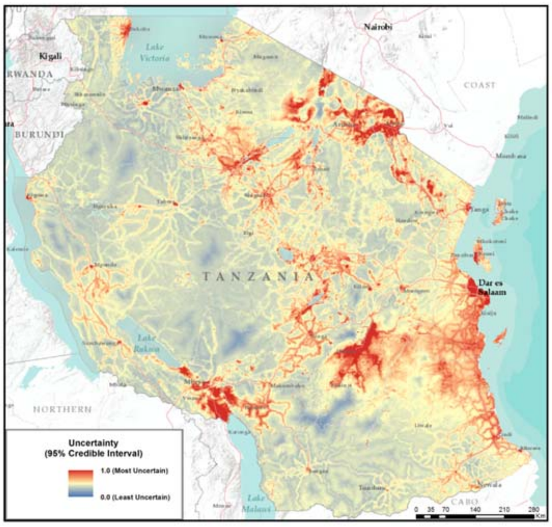

(d)

Figure 21. (a) Predicted map of the below $1.25 a day headcount ratio for Tanzania. This displays the mean value of the predictive posterior distribution at each 1x1km pixel. Major waterbodies and city names are also overlaid for context. (b) The precision of the <$1.25 a day model output for Tanzania as measured using the 95% credible interval. Major waterbodies and city names are also overlaid for context. (c) Predicted map of the below $2 a day headcount ratio for Tanzania. This displays the mean value of the predictive posterior distribution at each 1x1km pixel. Major waterbodies and city names are also overlaid for context. (d) The precision of the <$2 a day model output for Tanzania as measured using the 95% credible interval. Major waterbodies and city names are also overlaid for context.

6.0 GRIDDED POPULATION MAPS AND ADDITIONAL DATASETS

The poverty maps constructed here provide estimates of the proportion of the population in each 1x1km grid square that are classed as living in poverty by the MPI or consumption-based metrics. To convert these to estimates of absolute numbers of people living in poverty, estimates of the number of people residing in these grid cells are required. These were produced through the WorldPop project (www.worldpop.org.uk), and updated with the most recent and spatially detailed census data as part of this project. Full details of the input datasets and mapping methodologies used can be found on the Worldpop project website and the papers linked from the sites. Figures 22, 23 and 24 show the output population datasets for the six study countries.

As additional deliverables, spatial datasets on road networks (derived from a combination of national agency datasets and gRoads (http://www.ciesin.columbia.edu/confluence/display/roads/Global+Roads+Data), accessibility (as in figure 9), satellite nightlights (as in figure 7), livestock (from UN FAO: http://www.fao.org/ag/againfo/resources/en/glw/home.html), and age structure proportions (where available through WorldPop) for study countries were also constructed.

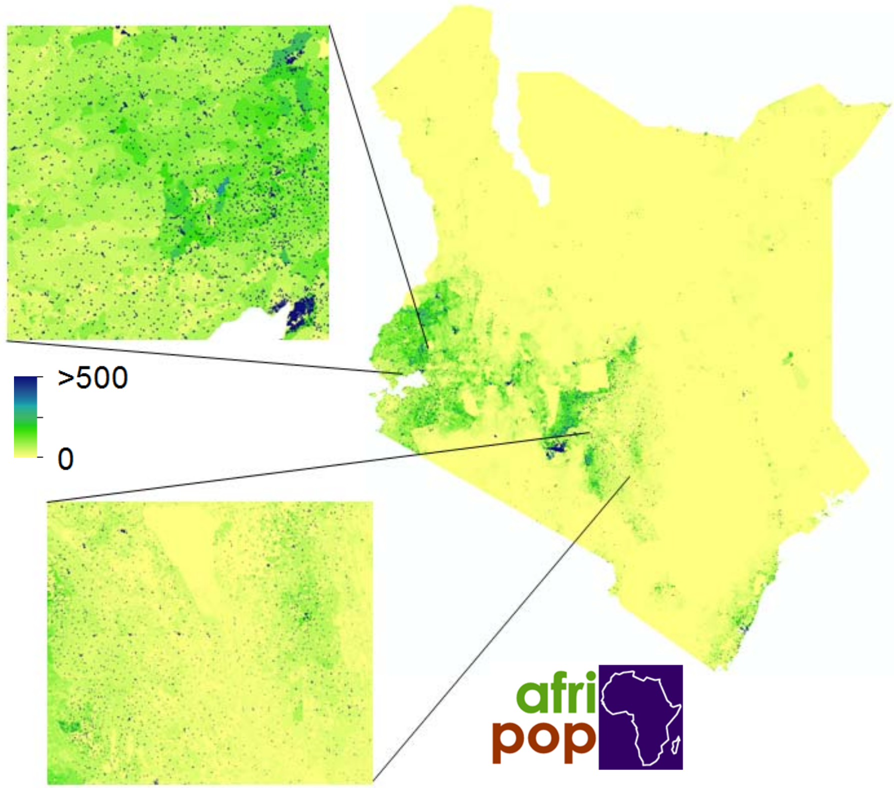

(a)

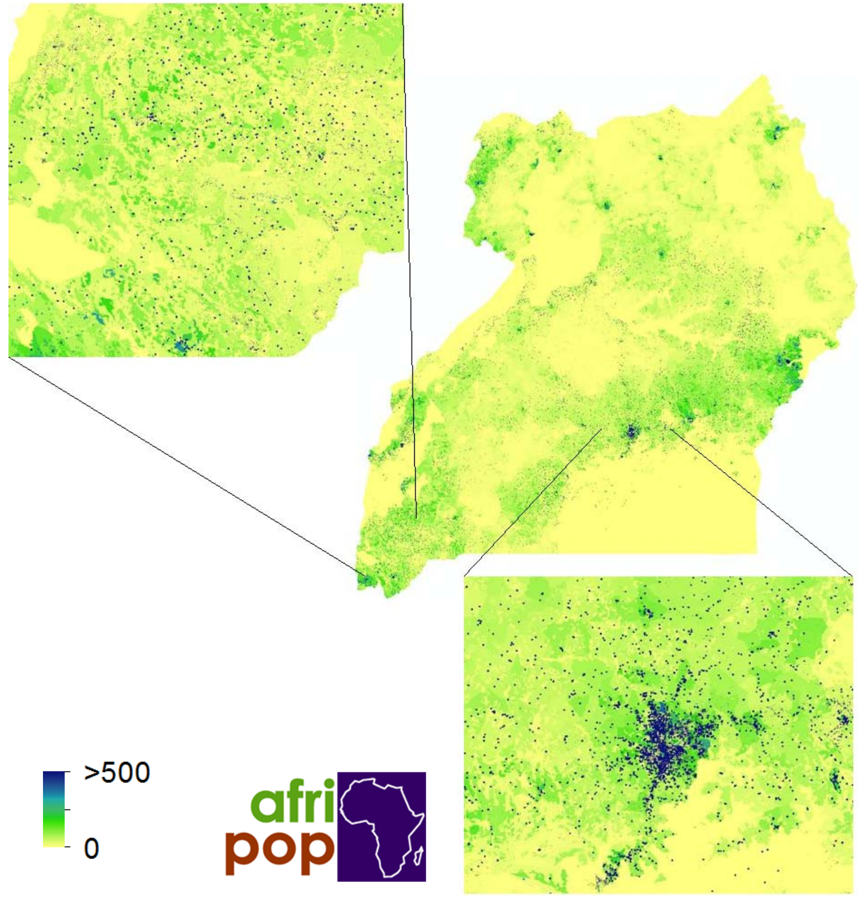

(b)

Figure 22. Estimated number of people residing in each 100x100m grid cell in 2010 for (a) Kenya and (b) Uganda. The inset boxes show close-ups of specific areas.

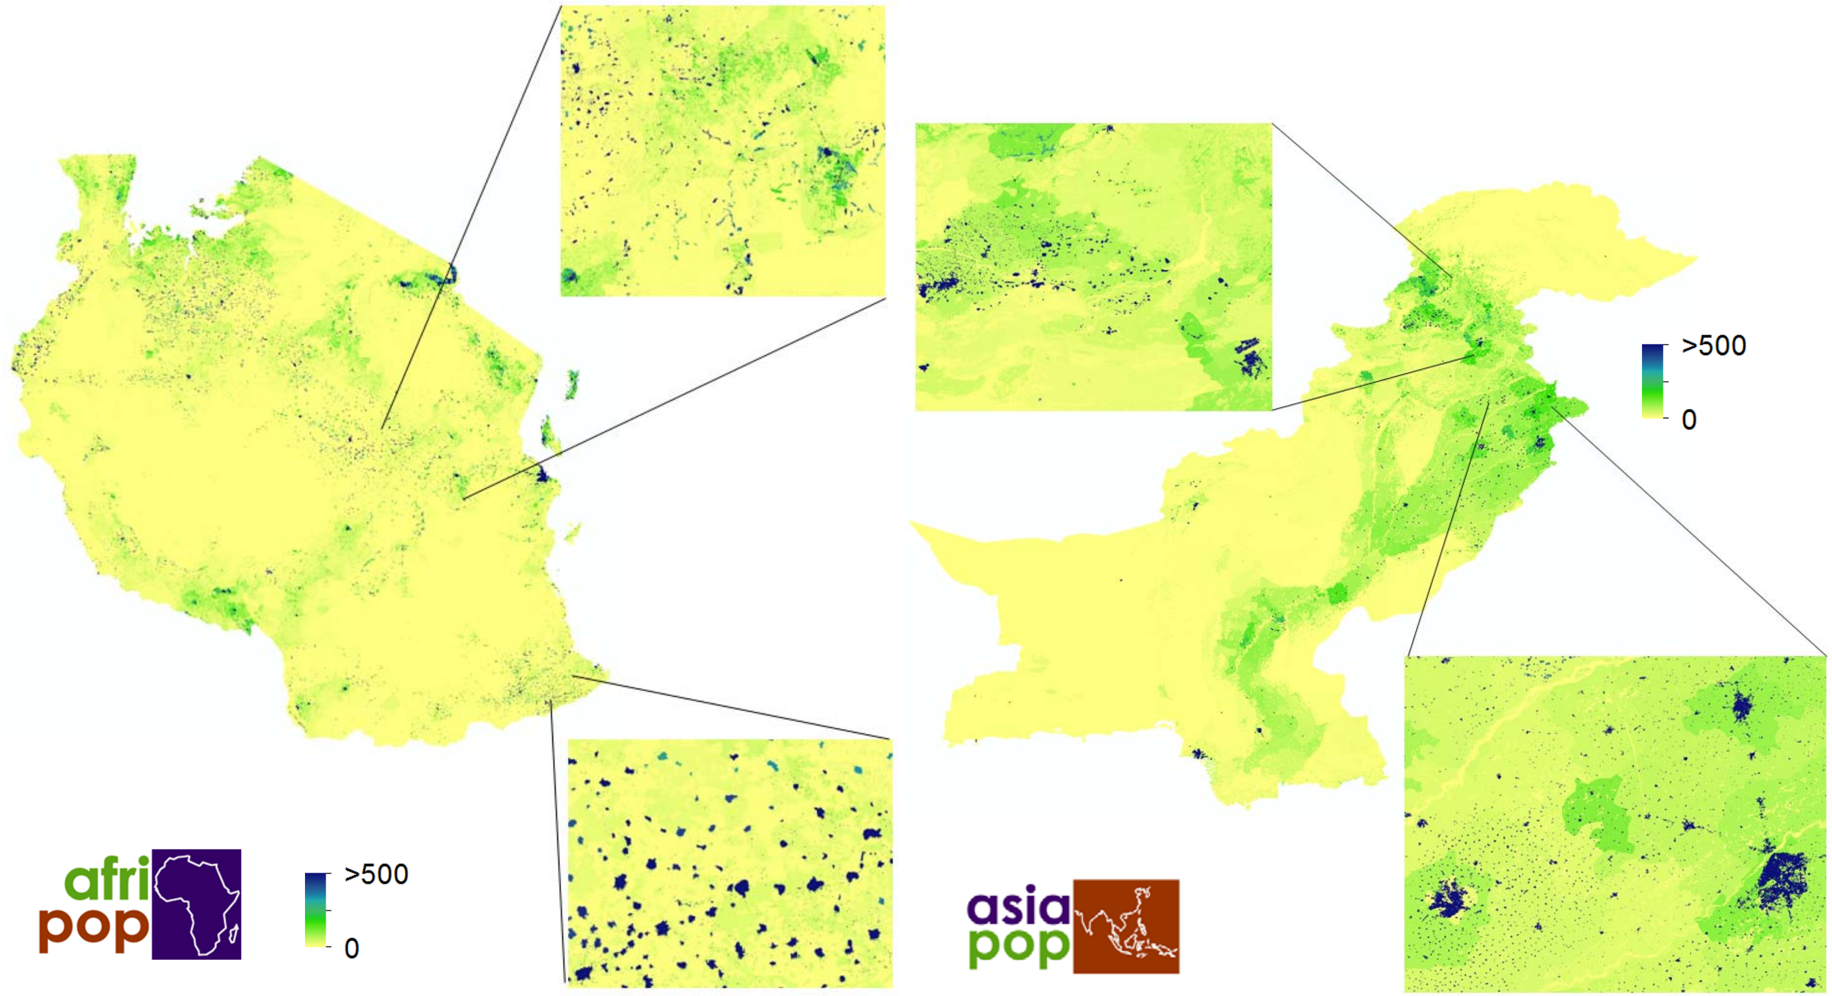

(a) (b)

Figure 23. Estimated number of people residing in each 100x100m grid cell in 2010 for (a) Tanzania and (b) Pakistan. The inset boxes show close-ups of specific areas.

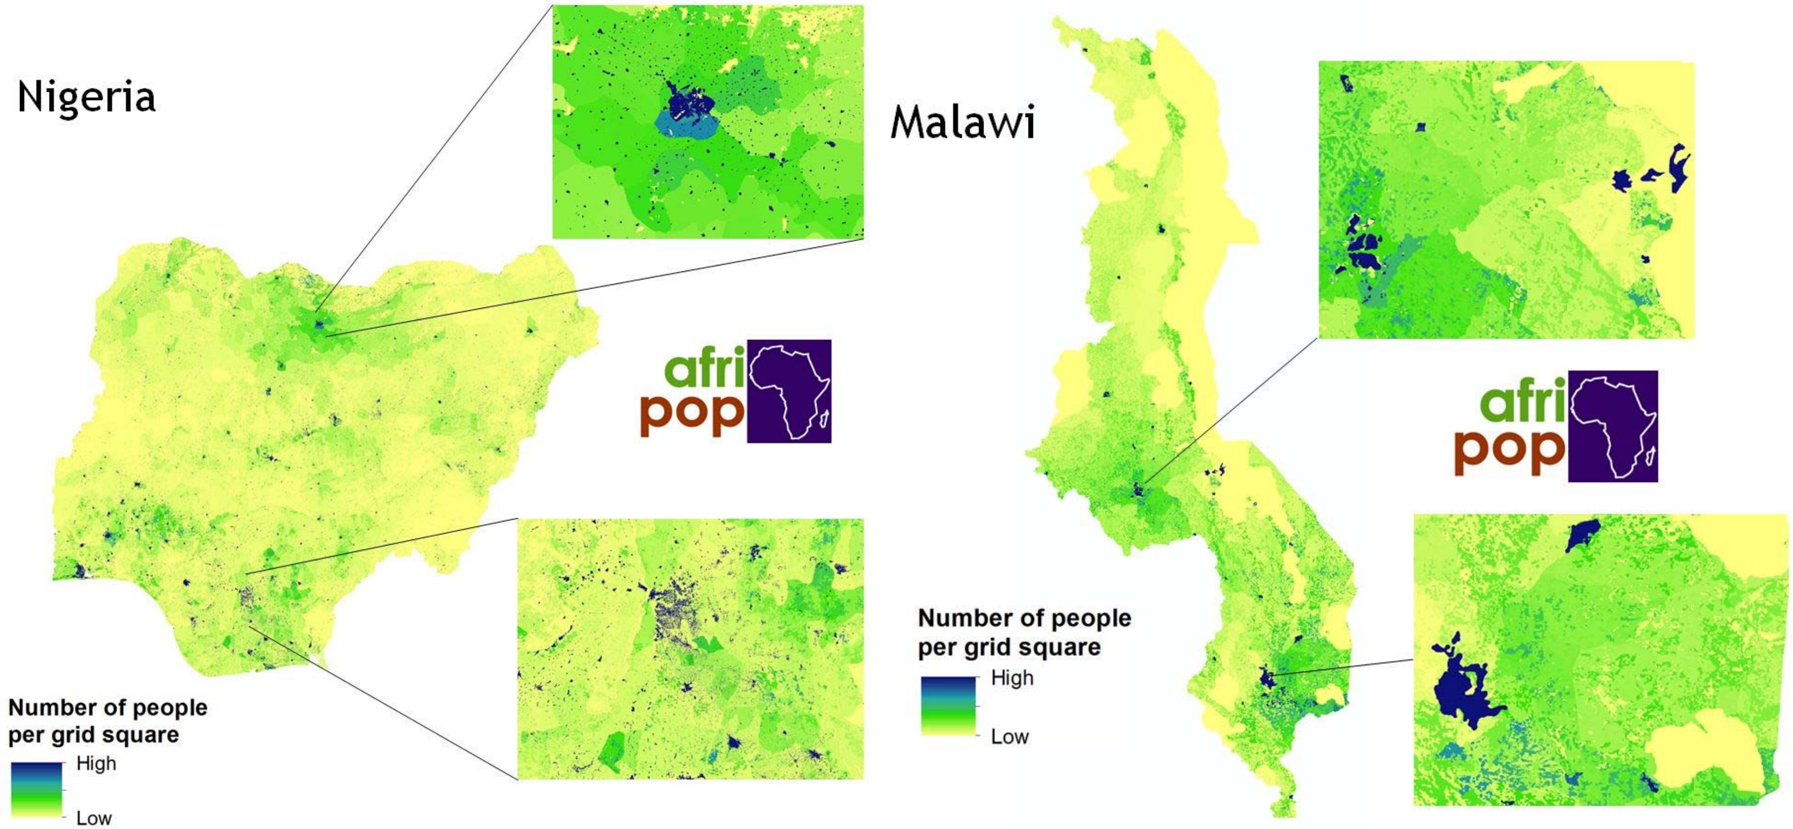

Figure 24. Estimated number of people residing in each 100x100m grid cell in 2010 for (a) Nigeria and (b) Malawi. The inset boxes show close-ups of specific areas.

7.0 NEXT STEPS

7.1 MAP COMPARISONS, SMALL AREA ESTIMATION HYBRID APPROACHES AND ADDITIONAL SURVEY DATASETS

Work is continuing in collaboration with Peter Lanjouw and his poverty mapping team at the World Bank. These collaborations include three sets of activities: (i) quantitative comparisons between our methodology presented here and the census-based World Bank poverty mapping methods, focused initially on Malawi, Nepal and Vietnam, where both census data and geolocated household surveys exist from approximately the same time period; (ii) the development of hybrid mapping approaches that draw on the strengths of the World Bank small area estimation methods and the Bayesian model based geostatistical methods presented here; (iii) discussions on intercomparisons/conversions between multidimensional poverty indices, which are more widely available through the DHS program, and consumption-based measures.

7.2 METHODOLOGICAL ADVANCEMENTS

We will continue to explore additional methodological advancements. The work presented here represents a first proof-of-concept exercise, with the exploration of many possible methodological improvements planned for future work. These will likely include: (i) exploring the possibility of a hierarchical modeling approach that implements high spatial resolution mapping in large cities (e.g. Nairobi and Mombasa in the case of Kenya, or Kampala in Uganda) to account for the significant variations in rates of poverty seen at small spatial scales in cities. Work has begun on assembling high resolution land use and satellite datasets for this work, but is currently limited by the offsets applied to the GPS coordinates in the DHS surveys that make them unsuitable for high resolution urban mapping. We anticipate that collaborations with the World Bank will provide us with access to LSMS data with non-offset GPS coordinates; (ii) developing approaches to account for the uncertainty introduced through these random geopositional offsets of cluster locations within DHS GPS survey data; (iii) the implementation of a space-time model – the work described so far has involved the implementation of a purely spatial model. With surveys being conducted across multiple years for many countries, the possibility exists to make use of such additional data from early surveys within a space-time framework to (i) improve model accuracies and (ii) enable predictions of poverty rates at multiple time periods.

7.3 ADDITIONAL VARIABLES AND COVARIATES

In addition to the methodological advancements outlined above, further steps in improving both the range of covariates used in the mapping, and the range of different variables to be mapped are planned: (i) An upcoming collaboration with Flowminder (www.flowminder.org) will explore the potential of spatial covariates representing mobility, consumption and social network structure derived from mobile phone call detail records in improving poverty mapping; (ii) An upcoming collaboration with the Bill and Melinda Gates Foundation polio team will explore the potential of the approaches developed here for demographic age structure mapping; (iii) Continued work through the Malaria Atlas Project will extend these approaches to new health-related indicator mapping; (iv) Continued collaborations with the Demographic and Health

Survey program will involve application of the developed approaches to DHS variables beyond the MPI analyses here.

8.0 REFERENCES

1. Tatem AJ, Adamo S, Bharti N, Burgert CR, Castro M, et al. (2012) Mapping populations at risk: improving spatial demographic data for infectious disease modeling and metric derivation. Popul Health Metr 10: 8.
2. Gething PW, Patil AP, Smith DL, Guerra CA, Elyazar IR, et al. (2011) A new world malaria map: Plasmodium falciparum endemicity in 2010. Malar J 10: 378.
3. Alkire S, Foster JE (2011) Understandings and Misunderstandings of Multidimensional Poverty Measurement Journal of Economic Inequality 9: 289-314.
4. Falkingham J, Namazie C (2002) Measuring Health and Poverty: A review of approaches to identifying the poor.
5. Alkire S, Foster J (2007) Counting and multidimensional poverty measures. Oxford Poverty and Human Development Initiative, Oxford. Working Paper Nr 7.
6. Alkire S, Santos M (2010) Acute Multidimensional Poverty: A New Index for Developing Countries. UNDP–HDRO, New York.
7. United Nations Development Program (2010) Human Development Report. UNDP, New York.
8. Uganda Bureau of Statistics (2010) Uganda National Household Survey Findings 2009/2010.
9. National Bureau of Statistics (NBS) [Tanzania] (2011) Tanzania National Panel Survey Report (NPS) - Wave 2, 2010 - 2011. Dar es Salaam, Tanzania: NBS.
10. Deaton A, Zaidi S (2002) Guidelines for Constructing Consumption Aggregates for Welfare Analysis: World Bank.
11. Coudouel A, Hentschel J, Wodon Q (2002) Poverty Measurement and Analysis. Poverty Reduction Strategy Paper (PRSP) Sourcebook. Washington D.C.: World Bank.
12. Montgomery MR (2003) Measuring living standards: household consumption and wealth indices. Washington, DC: The World Bank.
13. World Bank (2011) Defining Welfare Measures.
14. World Bank (2005) Measures of Poverty and Inequality Measures. WBI's Basic Poverty Measurement and Diagnostics course, Review poverty and inequality indicators: the World Bank Institute.
15. Samuelson PA (1974) Complementarity-An Essay on the 40th Anniversary of the Hicks-Allen Revolution in Demand Theory. Journal of Economic Literature, American Economic Association 12: 1255-1289.
16. Bellu' LG, Liberati P (2005) Equivalence Scales General Aspects. Food and Agriculture Organization of the United Nations, FAO.
17. Ravallion M, Chen S, Sangraula P (2008) Dollar a Day Revisited. World Bank Publications.
18. Development Research Group of the World Bank (2014) PovcalNet: an online poverty analysis tool.
19. World Bank (2014) Research, PovcalNet: an online poverty analysis tool, Methodology, The underlying data.
20. Sillers D, U.S. Agency for International Development Calculating PPP Conversion Factors and “$1-a-day” Poverty Lines. From the Annex to National and International Poverty Lines: An Overview.
21. World Bank (2014) Data, Indicators, Consumer price index.
22. Malawi Government - National Statistical Office (March 2012) Malawi Third Integrated Household Survey (IHS3) 2010-2011 - Basic Information Document.

23. Diggle PJ, Ribeiro PJ (2007) Model-based geostatistics; Bickel P, Diggle P, Fienberg S, Gather U, Olkin I et al., editors. New York: Springer. 228 p.
24. Diggle PJ, Tawn JA, Moyeed RA (1998) Model-based geostatistics. Journal of the Royal Statistical Society Series C-Applied Statistics 47: 299-326.
25. Rue H, Martino S, Chopin N (2009) Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. Journal of the Royal Statistical Society Series B: Statistical Methods 71: 319-392.
26. R Development Core Team (2008) R: a language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing, URL: http://www.R-project.org.
27. Antosiewicz HA (1964) Bessel Functions of Integer Order. In: Abramowitz M, Stegun IA, editors. Handbook of Mathematical Functions: with Formulas, Graphs, and Mathematical Tables. New York, U.S.A.: Dover Publications Inc. pp. 435-478.
28. Davis GM (1964) Gamma function and related functions. In: Abramowitz M, Stegun IA, editors. Handbook of Mathematical Functions: with Formulas, Graphs, and Mathematical Tables. New York, U.S.A.: Dover Publications Inc. pp. 253-293.

9.0 ACKNOWLEDGEMENTS

We are grateful to the following people for their advice and input: Karina Nielsen and Jake Kendall at the Bill and Melinda Gates Foundation, Clara Burgert and her team at Measure DHS, Peter Lanjouw and his team at the World Bank, Catherine Linard, Nirav Patel, Andrea Gaughan and Forrest Stevens at WorldPop.