Poverty Assessment Tool Accuracy Submission
USAID/IRIS Tool for Nigeria
Submitted: August 5, 2011

The following report is divided into five sections. Section 1 describes the data used to create the Poverty Assessment Tool for Nigeria. Section 2 details the set of statistical procedures used for selecting indicators and for estimating household expenditure or, for some models, the probability that a household is very poor. Section 3 reports on the in-sample accuracy of each prediction model considered. Sections 4 and 5 explain how regression coefficients are used in poverty prediction and how these predictions are used to classify households into the "very poor" and "not very poor" categories.

Annex 1 to this report provides accuracy results for an additional poverty line beyond that required by the Congressional legislation. Annex 2 reviews the out-of-sample accuracy for the Nigeria Poverty Assessment Tool.

1. Data source

For Nigeria, existing data from the 2003/2004 Nigeria Living Standard Survey (NLSS) were used to construct the poverty assessment tool. The full sample of 19,158 households is nationally representative. The sample used for tool construction comprises a randomly selected 14,369 households (75 percent of the full sample). The remainder, another randomly selected 4,789 households, is reserved for out-of-sample accuracy testing, which will investigate the robustness of in-sample poverty estimation.

2. Process used to select included indicators

Suitable household surveys, such as the LSMS, typically include variables related to education, housing characteristics, consumer durables, agricultural assets, and employment. Items included in the consumer durables and agricultural assets modules of the survey are commonly used as variables in PAT analyses. Due to the large number of, what appeared to be, nonrandom missing values in these sections of the survey results, respondent asset and durable ownership patterns were not included in this analysis. For Nigeria, more than 42 indicators from all categories were considered.

The MAXR procedure in SAS was used to select the best poverty indicators (for variables found to be practical) from the pool of potential indicators in an automated manner. MAXR is commonly used to narrow a large pool of possible indicators into a more limited, yet statistically powerful, set of indicators. The MAXR technique seeks to maximize explained variance (i.e., R^2) by adding one variable at a time (per step) to the regression model, and then considering all combinations among pairs of regressors to move from one step to the next. Thus, the MAXR technique allows us to identify the best model containing 15 variables (not including control variables for household size, age of the household head, and location).

The MAXR procedure yielded the best 15 variables for the OLS model (also used for the Quantile model) and another set of the best 15 variables for the Linear Probability model (also used for the Probit model). The final set of indicators and their weights, therefore, depended on selecting one of these four statistical models—OLS, Quantile, Linear Probability, or Probit—as the best model.[1] This selection of the best model was based on the Balanced Poverty Accuracy Criterion (BPAC) and the Poverty Incidence Error (PIE), along with practicality considerations.[2]

3. Estimation methods used to identify final indicators and their weights/coefficients

As explained more fully in Section 5, the line used to construct the poverty tool for Nigeria is the $1.25/day line. Table 1 summarizes the accuracy results achieved by each of the eight estimation methods in predicting household poverty relative to this poverty line. For Nigeria, the most accurate method, on the basis of BPAC, is the 2-step Quantile regression. However, the 1-step Quantile regression is only slightly less accurate and requires only 15 indicators. Following precedent from previous decisions made in consultation with USAID, the 1-step Quantile was selected as the best model, taking into consideration both accuracy and practicality.

Table 1: In-sample Accuracy Results for Prediction at the Legislative Poverty Line

NIGERIA ($1.25/day line*) Share of "Very Poor": 58.2%	**Total Accuracy**	**Poverty Accuracy**	**Under-coverage**	**Leakage**	**PIE**	**BPAC**
Single-step methods						
OLS	73.18	82.88	17.12	28.70	6.78	71.30
Quantile regression (estimation point: 54 percentile)	**72.65**	**76.52**	**23.48**	**23.25**	**-0.14**	**76.29**
Linear Probability	73.09	81.28	18.72	27.24	4.99	72.76
Probit	73.02	80.97	19.03	27.06	4.70	72.94
Two-step methods						
OLS – 59 percentile cutoff	73.16	82.46	17.54	28.32	6.31	71.68
Quantile (estimation points: 45, 27) 59 percentile cutoff	72.73	76.77	23.23	23.37	0.08	76.63
LP – 85 percentile cutoff	73.26	81.38	18.62	27.07	4.94	72.93
Probit –85 percentile cutoff	73.09	81.01	18.99	26.99	4.68	73.01
*The $1.25 per day per capita international poverty line in 2005 Purchasing Power Parity terms is 81.46 Nigerian Naira in January 2004 prices.						

[1] The set of indicators and their weights also depended on the selection of a 1-step or 2-step statistical model.

[2] For a detailed discussion of these accuracy criteria, see "Note on Assessment and Improvement of Tool Accuracy" at www.povertytools.org.

For Nigeria, the functionality of predicting the poverty rate at other poverty lines—in this case, the $0.75/day, $1.00/day, $2.00/day, and $2.50/day —has been added. This functionality is based on statistical models for prediction at the $1.25/day and $2.50/day lines. The methodology and the accuracy results for this prediction are discussed in Annex 1.

4. How coefficients and weights are used to estimate poverty status or household expenditures

For the quantile regression method, the estimated regression coefficients indicate the weight placed on each of the included indicators in estimating the household expenditures of each household in the sample. These estimated coefficients are shown in Table 3. In constructing the Poverty Assessment Tool for each country, these weights are inserted into the "back-end" analysis program of the CSPro template used to calculate the incidence of extreme poverty among each implementing organization's clients.

5. Decision rule used for classifying households as very poor and not very-poor

The legislation governing the development of USAID tools defines the "very poor" as either (1) those living on the local equivalent of less than the international poverty line ($1.25/day in 2005 PPP terms)[3] or (2) the poorest half of those living below the poverty line established by the national government. The applicable poverty line for developing USAID tools is the one that yields the higher household poverty rate in a given country.

In Nigeria the applicable threshold is the international poverty line of $1.25/day, at the level of prices prevailing when the household survey data were collected (2003-2004). The value of the line in those prices is 29753.17 Nigerian Naira per capita per year.[4] At these values, the $1.25/day poverty line identifies 58.2% of households as "very poor" and 64.2% of the population as "very poor."

By comparison, in 2003-2004 45.1% of Nigerian households lived below the national poverty line of 2,511 Nigerian Naira per adult equivalent per month. The second definition of extreme poverty in the legislation identifies 22.6% of households as "very poor."

Hence the decision rule for Nigeria's USAID poverty assessment tool in classifying the "very poor" (and the "not very-poor") is whether that predicted per capita daily expenditures of a household fall below (or above) the $1.25/day poverty line.

---

[3] The legislation specifies the international poverty line as the "equivalent of $1 per day (as calculated using the purchasing power parity (PPP) exchange rate method)." USAID and IRIS interpret this to mean the international poverty line used by the World Bank to track global progress toward the Millennium Development Goal of cutting the prevalence of extreme poverty in half by 2015. The World Bank adopted the $1.25/day line in 2008 to incorporate improved estimates of PPP based on data from 2005. The applicable 2005 PPP rate for Nigeria is 78.58.

[4] The calculation for the $1.25/day poverty line is (1.25*78.582528584489)*(365.25)*(82.9291/100) where the final term is the CPI adjustment from average 2005 prices to prices at the time of the survey.

Because the selected tool is based on a Quantile model, each household whose estimated per capita consumption expenditures according to the tool is less than or equal to the $1.25/day poverty line is identified as “very poor,” and each household whose estimated per capita consumption expenditures exceeds the $1.25/day poverty line is identified as “not very-poor.”

Table 2 below compares the poverty status of the sample households as identified by the selected model, versus their true poverty status as revealed by the data from the benchmark household survey (in-sample test). The upper-left and lower-right cells show the number of households correctly identified as “very poor” or “not very-poor,” respectively. Meanwhile, the upper-right and lower-left cells indicate the twin errors possible in poverty assessment: misclassifying very poor households as not very-poor; and the opposite, misclassifying not very-poor households as very poor.

Table 2: Poverty Status of Sample Households, as Estimated by Model and Revealed by the Benchmark Survey

	Number of households identified as very-poor by the tool	**Number of households identified as not very-poor by the tool**
Number of “true” very-poor households (as determined by benchmark survey)	**6,436** (44.8%)	**1,975** (13.7%)
Number of “true” not very-poor households (as determined by benchmark survey)	**1,955** (13.6%)	**4,003** (27.9%)

Table 3: Regression Estimates using 1-step Quantile Method for Prediction at the $1.25/day Poverty Line

.54 Quantile regression — Number of obs = 14,369
Min sum of deviations = 7265.404 — Pseudo R2 = 0.1962

Variable	Coef.	Std. Err.	t	P>|t|	[95% Conf. Interval]	
Intercept	10.9117	0.0815	133.9500	0.0000	10.7520	11.0714
Household size	-0.2437	0.0087	-28.0200	0.0000	-0.2608	-0.2267
Household size squared	0.0106	0.0006	19.1500	0.0000	0.0095	0.0117
Household head age	0.0039	0.0031	1.2500	0.2100	-0.0022	0.0100
Household head age squared	0.0000	0.0000	-1.3000	0.1950	-0.0001	0.0000
Household lives in rural area	-0.1890	0.0231	-8.1800	0.0000	-0.2343	-0.1437
Household lives in South Central	-0.1281	0.0266	-4.8200	0.0000	-0.1802	-0.0760
Household lives in North East	-0.0684	0.0253	-2.7000	0.0070	-0.1180	-0.0188
Household lives in South	-0.1187	0.0315	-3.7700	0.0000	-0.1803	-0.0570
Household lives in the South West	-0.0661	0.0306	-2.1600	0.0310	-0.1260	-0.0062
Household lives in South East	0.1731	0.0310	5.5800	0.0000	0.1123	0.2339
Share of members (excluding head) that never achieved an education qualification	-0.2703	0.0379	-7.1300	0.0000	-0.3446	-0.1960
Share of members (excluding head) whose highest education qualification is the First School Leaving Certificate	-0.1499	0.0538	-2.7900	0.0050	-0.2552	-0.0445
Share of members (excluding head) whose highest education qualification is the National Certificate of Education or National Diploma	0.5255	0.2082	2.5200	0.0120	0.1175	0.9335
Share of members (excluding head) whose highest education qualification is completion of a bachelors or a higher national diploma program	0.9548	0.2019	4.7300	0.0000	0.5591	1.3506
Household head's highest education qualification is the First School Leaving Certificate	-0.0600	0.0219	-2.7400	0.0060	-0.1030	-0.0171
Household head can read a letter in English	0.1587	0.0284	5.5800	0.0000	0.1029	0.2144
Household head is able to do written calculations	0.0703	0.0274	2.5600	0.0100	0.0165	0.1241
Household head involved in farming, livestock or fishing in last 12 months	0.0623	0.0222	2.8000	0.0050	0.0187	0.1059

Variable	Coef.	Std. Err.	t	P>|t|	[95% Conf. Interval]	
Number of rooms in dwelling	0.0434	0.0054	8.0500	0.0000	0.0328	0.0539
Household shares its dwelling with other household(s)	-0.0516	0.0197	-2.6100	0.0090	-0.0902	-0.0129
Main source of drinking water from borehole/hand pump	-0.0472	0.0210	-2.2500	0.0250	-0.0883	-0.0060
Main lighting source is electricity from a central source	0.1134	0.0211	5.3900	0.0000	0.0721	0.1546
Household's toilet is on water	0.1611	0.0394	4.0900	0.0000	0.0838	0.2384
Floor is made of earth or mud	-0.1049	0.0191	-5.5000	0.0000	-0.1423	-0.0675
Roof is made of mud or mud bricks	-0.1077	0.0256	-4.2000	0.0000	-0.1579	-0.0574

Annex 1: Poverty Prediction at the $2.50/day Poverty Line and Discussion of Additional Poverty Lines

Strictly construed, the legislation behind the USAID poverty assessment tools concerns "very poor" and "not very-poor" beneficiaries. Nevertheless, the intended outcome of the legislation is to provide USAID and its implementing partners with poverty measurement tools that they will find useful.

After discussions among USAID, IRIS, and other members of the microenterprise community, a consensus emerged that the tools would benefit from predictive capacity beyond legislatively-defined extreme poverty. To that end, on agreement with USAID, IRIS has used the best indicators and regression type for predicting the "very poor" to also identify the "poor." For $1.25/day PPP models, this will be the $2.50/day PPP; for median poverty models, the "poor" threshold will be the national poverty line. Following this logic, then, the "poor" ("not poor") in Nigeria are defined as those whose predicted expenditures fall below (above) the $2.50/day poverty line.

Table 4 summarizes the predictive accuracy results for the $2.50/day poverty line using the Quantile model specification from the $1.25/day poverty line. The indicators are the same as those in the model for the $1.25/day line, but the percentile of estimation and the coefficients of the model were allowed to change (compare Tables 3 and 6). This methodology allows the content and length of the questionnaire to remain the same, but permits greater accuracy in predicting at the $2.50/day poverty line.

Based on the statistical models underlying prediction at these two lines, IRIS has also introduced the functionality of prediction at five lines to increase the usefulness of the tool to partner organizations. For Nigeria, these five lines are the $0.75/day line, $1.00/day line, $1.25/day line, $2.00/day line, and the $2.50/day line. Poverty rates at the first three lines are predicted using the best model for the $1.25/day line, while poverty rates at the last two lines are predicted using the best model for the $2.50/day line. As discussed in this document, accuracy has been tested at the $1.25 and $2.50 lines. Given this, the predictions made at the other lines are intended for indicative use by implementing partners.

The tabulation of poverty prevalence has also been expanded to provide a fuller summary of the incidence of poverty among the implementing organization's clients. Poverty status at the five poverty lines is cross tabulated with regional location, household head's gender, household head's education by gender, household size, and housing conditions. Again, the additional information provided is for indicative purposes rather than statistical inference.

Table 4: Accuracy Results Obtained for Prediction at the $2.50/day Poverty Line

NIGERIA $2.50/day Line Share of Poor: 69%	**Total Accuracy**	**Poverty Accuracy**	**Under-coverage**	**Leakage**	**PIE**	**BPAC**
Single-step methods						
Quantile regression (estimation point: 69 percentile)	85.18	91.53	8.47	8.68	0.18	91.32

Table 5 below compares the poverty status of the sample households as identified by the selected model, versus their true poverty status as revealed by the data from the benchmark household survey (in-sample test). The upper-left and lower-right cells show the number of households correctly identified as "poor" or "not poor," respectively. Meanwhile, the upper-right and lower-left cells indicate the twin errors possible in poverty assessment: misclassifying poor households as not poor; and the opposite, misclassifying not poor households as poor.

Table 5: Poverty Status of Sample Households, as Estimated by Model and Revealed by the Benchmark Survey, at $2.50/day Line

	Number of households identified as very-poor by the tool	**Number of households identified as not very-poor by the tool**
Number of "true" very-poor households (as determined by benchmark survey)	**11,366** (79.1%)	**1,052** (7.3%)
Number of "true" not very-poor households (as determined by benchmark survey)	**1,077** (7.5%)	**874** (6.1%)

Table 6: Regression Estimates using 1-Step Quantile Method for Prediction at $2.50/day Poverty Line

.69 Quantile regression   Number of obs = 14,369
Min sum of deviations = 6403.131   Pseudo R2 = 0.2048

Variable	Coef.	Std. Err.	t	P>|t|	[95% Conf. Interval]	
Intercept	11.1092	0.0749	148.3200	0.0000	10.9624	11.2560
Household size	-0.2344	0.0078	-29.9600	0.0000	-0.2497	-0.2190
Household size squared	0.0104	0.0005	21.8600	0.0000	0.0095	0.0114
Household head age	0.0023	0.0029	0.7900	0.4320	-0.0034	0.0079
Household head age squared	0.0000	0.0000	-0.6000	0.5510	-0.0001	0.0000
Household lives in rural area	-0.1925	0.0214	-9.0000	0.0000	-0.2344	-0.1506
Household lives in South Central	-0.0684	0.0253	-2.7000	0.0070	-0.1179	-0.0188
Household lives in North East	-0.0640	0.0238	-2.6800	0.0070	-0.1107	-0.0172
Household lives in South South	-0.0723	0.0297	-2.4300	0.0150	-0.1306	-0.0140
Household lives in the South West	-0.0487	0.0290	-1.6800	0.0930	-0.1055	0.0081
Household lives in South East	0.1697	0.0292	5.8200	0.0000	0.1125	0.2269
Percent of family members that never achieved an education qualification	-0.2953	0.0350	-8.4300	0.0000	-0.3639	-0.2266
Percent of family members whose highest education qualification is the First School Leaving Certificate	-0.1898	0.0497	-3.8200	0.0000	-0.2872	-0.0923
Percent of family members whose highest education qualification is the National Certificate of Education or National Diploma	0.6906	0.1947	3.5500	0.0000	0.3089	1.0723
Percent of family members whose highest education qualification is completion of a bachelors or a higher national diploma program	1.0598	0.1903	5.5700	0.0000	0.6868	1.4328
Household head's highest education qualification is the First School Leaving Certificate	-0.0524	0.0203	-2.5800	0.0100	-0.0922	-0.0126

Household head can read a letter in English	0.1509	0.0263	5.7300	0.0000	0.0993	0.2025
Household head is able to do written calculations	0.1041	0.0257	4.0500	0.0000	0.0537	0.1546
Household head involved in farming, livestock or fishing in last 12 months	0.0456	0.0204	2.2400	0.0250	0.0057	0.0856
Number of rooms in dwelling	0.0452	0.0049	9.2500	0.0000	0.0356	0.0547
Household shares its dwelling with other household(s)	-0.0444	0.0183	-2.4300	0.0150	-0.0802	-0.0086
Main source of drinking water from borehole/handpump	-0.0404	0.0194	-2.0800	0.0380	-0.0785	-0.0023
Main lighting source is electricity from a central source	0.1132	0.0192	5.8900	0.0000	0.0755	0.1508
Household's toilet is on water	0.1581	0.0354	4.4600	0.0000	0.0887	0.2275
Floor is made of earth or mud	-0.1227	0.0175	-7.0000	0.0000	-0.1571	-0.0883
Roof is made of mud or mud bricks	-0.1061	0.0238	-4.4600	0.0000	-0.1528	-0.0595

Annex 2: Out-of-Sample Accuracy Tests

In statistics, prediction accuracy can be measured in two fundamental ways: with in-sample methods and with out-of-sample methods. In the in-sample method, a single data set is used. This single data set supplies the basis for both model calibration and for the measurement of model accuracy. In the out-of-sample method, at least two data sets are utilized. The first data set is used to calibrate the predictive model. The second data set tests the accuracy of these calibrations in predicting values for previously unobserved cases.

The previous sections of this report provide accuracy results of the first type only. The following section presents accuracy findings of the second type, as both a supplement to certification requirements and as an exploration of the robustness of the best model outside of the 'laboratory' setting.

As noted in section 1, the data set used to construct the Nigeria tool was divided randomly into two data sets 14,369 households (75 percent of the sample) and 4,789 households (25 percent sample). A naïve method for testing out-of-sample accuracy—or for overfitting—is to simply apply the model calibrated on the first data set to the observations contained in the holdout data set. These results are show in Table 7. The best model (1-step quantile) performs well in terms of BPAC and PIE, lost 0.86 BPAC points and 0.52 points for PIE, respectively.

Table 7: Comparison of In-Sample and Out-of-Sample Accuracy Results

	Total Accuracy	Poverty Accuracy	Under-coverage	Leakage	PIE	BPAC
In-Sample Prediction						
	72.65	76.52	23.48	23.25	-0.14	76.29
Out-of-Sample Prediction						
	72.54	76.58	23.42	24.57	0.66	75.43

Another, more rigorous method for testing the out-of-sample accuracy performance of the tool is to provide confidence intervals for the accuracy measures, derived from 1,000 bootstrapped samples from the holdout sample.[5] Each bootstrapped sample is constructed by drawing observations, with replacement, from the holdout sample. The calibrated model is then applied to each sample to yield poverty predictions; across 1,000 samples, this method provides the sampling distributions for the model's accuracy measures.

Table 8 presents the out-of-sample, bootstrapped confidence intervals for the 1-step Quantile model. The performance of this model is very good. The confidence interval around the sample mean BPAC is relatively narrow at +/- 2.24 percentage points. For

[5] This method of out-of-sample testing is used by Mark Schreiner for the PPI scorecards as detailed on www.microfinance.com

PIE, which measures the difference between the predicted poverty rate and the actual poverty rate, the confidence interval is +/- 1.89 percentage points.

Table 8: Bootstrapped Confidence Intervals on Assumption of Normality

Accuracy Measure	Mean	Std. Dev.	Confidence interval	
			LB	UB
Total Accuracy	72.53	0.77	71.01	74.05
Poverty Accuracy	76.60	1.01	74.61	78.59
Undercoverage	23.40	1.01	21.41	25.39
Leakage	24.61	1.23	22.20	27.01
PIE	0.68	0.96	-1.20	2.57
BPAC	74.92	1.14	72.69	77.16

The results presented in Table 8 assume a normal distribution for the accuracy measures from the bootstrapped samples. This ignores the possibility that these estimates may have a skewed distribution. Table 9 presents alternative 95% confidence intervals. The lower bound is defined by the 2.5^{th} percentile of the sample distribution for each measure; the upper bound is defined by the 97.5^{th} percentile. On the whole, the results are quite similar between Tables 8 and 9.

Table 9: Bootstrapped Confidence Intervals Computed Empirically from Sampling Distribution without Normality Assumption

Accuracy Measure	Confidence interval	
	LB	UB
Total Accuracy	71.11	74.12
Poverty Accuracy	74.55	78.58
Undercoverage	21.42	25.45
Leakage	22.26	27.18
PIE	-1.14	2.53
BPAC	72.47	76.86

The primary purpose of the PAT is to assess the overall extreme poverty rate across a group of households. The out-of-sample results for PIE in Table 8 and Table 9 indicate that the extreme poverty rate estimate produced by the Nigeria PAT appears to be slightly biased toward overestimating the actual extreme poverty rate, but nonetheless has a relatively narrow confidence interval. By this measure, the predictive model behind the Nigeria PAT is accurate.