Poverty Assessment Tool Accuracy Submission
USAID/IRIS Tool for Ghana (Updated with new data)
Submitted: June 4, 2010

The following report is divided into five sections. Section 1 describes the data set used to create the Poverty Assessment Tool for Ghana. 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 Ghana Poverty Assessment Tool.

1. Data source

For Ghana, existing data from the 2005/6 Ghana Living Standards Survey (GLSS) were used to construct the poverty assessment tool. The full sample of 8,687 households is nationally representative. The sample used for tool construction comprises a randomly selected 4,344 households (50 percent of the full sample). The remainder, another randomly selected 4,343 households, is reserved for out-of-sample accuracy testing, which investigates the robustness of in-sample poverty estimation.

The previous PAT for Ghana was based on a 2001/2002 GLSS.

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, illness and disability, and employment. For Ghana, more than 75 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 Balance 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 Ghana is the international poverty line: $1.25 per capita per day in 2005 PPP. Table 1 summarizes the accuracy results achieved by each of the eight estimation methods in predicting household poverty relative to this poverty line. For Ghana, 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

GHANA $1.25/day line***** Share of "very poor": 18.9%	**Total Accuracy**	**Poverty Accuracy**	**Under-coverage**	**Leakage**	**PIE**	**BPAC**
Single-step methods --						
OLS	86.09	52.08	47.92	26.50	-4.00	30.66
Quantile regression (estimation point: 41 percentile)	**85.66**	**62.02**	**37.98**	**38.72**	**0.14**	**61.28**
Linear Probability	86.50	39.80	60.20	12.02	-9.01	-8.37
Probit	86.93	50.76	49.24	20.67	-5.34	22.19
Two-step methods --						
OLS – 27 percentile cutoff	86.93	56.09	43.91	26.01	-3.35	38.19
Quantile (estimation points: 42, 14) 27 percentile cutoff	86.45	64.38	35.62	36.88	0.24	63.12
LP –29 percentile cutoff	87.56	56.91	43.09	23.47	-3.67	37.29
Probit – 29 percentile cutoff	87.27	54.04	45.96	22.16	-4.45	30.24
* $1.25/day poverty line is 188,749 Cedis per capita per month in 2005 prices. The international poverty line is based on World Bank's calculations and the recent 2005 PPP exchange rates.						

[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 Ghana, the functionality of predicting the poverty rate at other poverty lines—in this case, the \$0.75/day, \$1.00/day, \$1.25/day, \$2.00/day and \$2.50/day lines—have 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 the bottom (poorest) 50 percent of those living below the poverty line established by the national government or those living on the local equivalent of less than the international poverty line (\$1.25/day in 2005 PPP terms)[3]. The applicable poverty line for USAID tool development is the one that yields the higher household poverty rate for a given country.

In Ghana the applicable threshold is the international poverty line of \$1.25/day in 2005 PPP terms.[4] The value of this line at the time of the survey is 188,749 Cedis per capita per month. This line identifies 18.9% of households as "very poor."

The alternative possibility for the poverty line is the "median poverty line." Expressed in prices prevailing at the time of data collection, the median line is 187,120 Cedis per adult equivalent for Ghana. This line identifies 14.2% of households in the sample as very poor.

Hence the decision rule for Ghana's USAID poverty assessment tool in classifying the "very poor" (and the "not very-poor") is whether predicted per capita monthly

[3] The congressional 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. This poverty line has recently been recalculated by the Bank to accompany new, improved estimates of PPP. The applicable 2005 PPP rate for Ghana is 4475.822.

[4] The World Bank's PovcalNet estimates the poverty rate at 29.9% using the international \$1.25 PPP line and population weights. It appears that PovcalNet uses the 2006 MICS dataset for this estimation. We obtain a result of 28.8% in our own tests using the 2006 GLSS data and population weights.

expenditures of a household is less than or equal to (or above) the international poverty line.

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 international poverty line is identified as “very poor,” and each household whose estimated per capita consumption expenditures exceeds the international 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)	504 (11.6%)	308 (7.1%)
Number of “true” not very-poor households (as determined by benchmark survey)	314 (7.3%)	3217 (74.0%)

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

GHANA 1 STEP MAXR/QUANT: variables from MAXR/OLS 100 percentile model
Regression results, estimation point of 41 percentile

.41 Quantile regression — Number of obs = 4344
Min sum of deviations 1627.111 — Pseudo R2 = 0.4035

Variable	Coef.	Std. Err.	t	P>|t|	[95% Conf. Interval]	
Intercept	13.4002	0.0762	175.8600	0.0000	13.2508	13.5496
Household size	-0.2411	0.0105	-22.8700	0.0000	-0.2617	-0.2204
Household size squared	0.0095	0.0007	14.4000	0.0000	0.0082	0.0108
Household head age	0.0024	0.0032	0.7600	0.4450	-0.0038	0.0087
Household head age squared	-0.0001	0.0003	-1.8200	0.0680	-0.0001	0.0000
HH lives in western region	0.0187	0.0337	0.5600	0.5780	-0.0473	0.0847
HH lives in central region	0.0168	0.0358	0.4700	0.6380	-0.0533	0.0870
HH lives in greater accra region	0.0261	0.0313	0.8300	0.4050	-0.0354	0.0875
HH lives in volta region	-0.1257	0.0366	-3.4300	0.0010	-0.1975	-0.0539
HH lives in eastern region	0.0634	0.0331	1.9100	0.0560	-0.0015	0.1284
HH lives in brong-ahafo region	-0.0930	0.0341	-2.7300	0.0060	-0.1598	-0.0262
HH lives in northern region	-0.2454	0.0376	-6.5200	0.0000	-0.3192	-0.1717
HH lives in upper east region	-0.7122	0.0419	-17.0100	0.0000	-0.7943	-0.6301
HH lives in upper west region	-0.9788	0.0443	-22.0800	0.0000	-1.0657	-0.8919
Household lives in rural area	-0.1045	0.0228	-4.5900	0.0000	-0.1491	-0.0599
HH owns one or more sewing machines	0.0772	0.0221	3.5000	0.0000	0.0339	0.1205
HH owns one or more gas stoves	0.1569	0.0347	4.5300	0.0000	0.0890	0.2249
HH owns one or more fans	0.1085	0.0274	3.9600	0.0000	0.0548	0.1622
HH owns one or more video players	0.1838	0.0320	5.7400	0.0000	0.1211	0.2465
HH owns one or more TVs	0.1404	0.0266	5.2800	0.0000	0.0883	0.1925
HH owns one or more irons	0.1687	0.0281	6.0100	0.0000	0.1136	0.2237
HH member owns a house	0.1220	0.0209	5.8500	0.0000	0.0812	0.1629
HH member owns land	0.1136	0.0210	5.4200	0.0000	0.0725	0.1547
Number of rooms occupied by household	0.0812	0.0098	8.3000	0.0000	0.0620	0.1003
HH disposes of refuse by dumping elsewhere	-0.1195	0.0228	-5.2400	0.0000	-0.1643	-0.0748
HH disposes of refuse by burning	-0.1698	0.0385	-4.4100	0.0000	-0.2454	-0.0943
HH uses a flush toilet	0.1754	0.0357	4.9200	0.0000	0.1055	0.2453
Floor of dwelling is earth, mud, or mud bricks	-0.1019	0.0271	-3.7600	0.0000	-0.1550	-0.0487
Household head has attained no highest educational qualification	-0.0877	0.0210	-4.1800	0.0000	-0.1288	-0.0465
Share of HH members who have attained no highest educational qualification	-0.2166	0.0475	-4.5600	0.0000	-0.3099	-0.1234

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 Ghana are defined as those whose predicted incomes fall below (above) the $2.50 line in 2005 PPP.

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. These 5 lines are the $0.75/day, $1.00/day, $1.25/day, $2.00/day and $2.50/day lines. 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/line. As discussed in this document, accuracy has been tested at the $1.25/day line and again at $2.50/day line (with a model calibrated to that line, but using variables from the best $1.25/day line). 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 characteristics, 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

Ghana National Line Share of Poor: 50.5%	**Total Accuracy**	**Poverty Accuracy**	**Under-coverage**	**Leakage**	**PIE**	**BPAC**
Single-step methods						
Quantile regression (estimation point: 47)	81.27	82.08	17.92	18.60	0.35	81.40

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 Poverty Line

	Number of households identified as poor by the tool	**Number of households identified as not poor by the tool**
Number of “true” poor households (as determined by benchmark survey)	1829 (42.1%)	399 (9.2%)
Number of “true” not poor households (as determined by benchmark survey)	414 (9.5%)	1702 (39.2%)

Table 6: Regression Estimates using 1-step Quantile Method for Prediction at the $2.50 Poverty Line

GHANA 1 STEP MAXR/QUANT: variables from MAXR/OLS 100 percentile model
Regression results, estimation point of 47 percentile

.47 Quantile regression Number of obs = 4344
Min sum of deviations 1668.559 Pseudo R2 = 0.4036

Variable	Coef.	Std. Err.	t	P>|t|	[95% Conf. Interval]	
Intercept	13.4220	0.0933	143.8900	0.0000	13.2391	13.6049
Household size	-0.2493	0.0125	-19.9400	0.0000	-0.2738	-0.2247
Household size squared	0.0099	0.0008	13.1400	0.0000	0.0084	0.0114
Household head age	0.0056	0.0039	1.4400	0.1510	-0.0020	0.0131
Household head age squared	-0.0001	0.0000	-2.2800	0.0230	-0.0002	0.0000
HH lives in western region	0.0254	0.0411	0.6200	0.5370	-0.0552	0.1060
HH lives in central region	0.0101	0.0438	0.2300	0.8180	-0.0758	0.0960
HH lives in greater accra region	0.0422	0.0384	1.1000	0.2710	-0.0330	0.1175
HH lives in volta region	-0.1449	0.0446	-3.2500	0.0010	-0.2323	-0.0575
HH lives in eastern region	0.0560	0.0406	1.3800	0.1680	-0.0236	0.1355
HH lives in brong-ahafo region	-0.1120	0.0417	-2.6900	0.0070	-0.1936	-0.0303
HH lives in northern region	-0.2437	0.0464	-5.2500	0.0000	-0.3346	-0.1527
HH lives in upper east region	-0.7023	0.0515	-13.6200	0.0000	-0.8034	-0.6012
HH lives in upper west region	-0.9563	0.0541	-17.6700	0.0000	-1.0624	-0.8502
Household lives in rural area	-0.1201	0.0279	-4.3000	0.0000	-0.1748	-0.0653
HH owns one or more sewing machines	0.0827	0.0270	3.0600	0.0020	0.0297	0.1357
HH owns one or more gas stoves	0.1562	0.0421	3.7100	0.0000	0.0737	0.2388
HH owns one or more fans	0.0989	0.0334	2.9600	0.0030	0.0334	0.1644
HH owns one or more video players	0.1802	0.0391	4.6100	0.0000	0.1035	0.2569
HH owns one or more TVs	0.1437	0.0326	4.4000	0.0000	0.0797	0.2076
HH owns one or more irons	0.1485	0.0343	4.3300	0.0000	0.0812	0.2158
HH member owns a house	0.1181	0.0255	4.6400	0.0000	0.0682	0.1680
HH member owns land	0.1215	0.0257	4.7300	0.0000	0.0712	0.1718
Number of rooms occupied by household	0.0848	0.0119	7.1500	0.0000	0.0616	0.1080
HH disposes of refuse by dumping elsewhere	-0.1305	0.0279	-4.6900	0.0000	-0.1852	-0.0759
HH disposes of refuse by burning	-0.1392	0.0470	-2.9600	0.0030	-0.2313	-0.0471
HH uses a flush toilet	0.1618	0.0432	3.7400	0.0000	0.0770	0.2465
Floor of dwelling is earth, mud, or mud bricks	-0.0927	0.0330	-2.8100	0.0050	-0.1573	-0.0281
Household head has attained no highest educational qualification	-0.0774	0.0256	-3.0200	0.0030	-0.1277	-0.0271
Share of HH members who have attained no highest educational qualification	-0.2133	0.0574	-3.7200	0.0000	-0.3258	-0.1008

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 Ghana tool was divided randomly into two data sets of nearly equal size (4,344 and 4,343 households). 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, losing about 5 points and 1 point, 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						
	85.66	62.02	37.98	38.72	0.14	61.28
Out-of-Sample Prediction						
	85.59	60.00	39.99	35.99	-0.76	56.00

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

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

around the sample mean BPAC is relatively narrow at +/- 7.8 percentage points. For PIE, which measures the difference between the predicted poverty rate and the actual poverty rate, the confidence interval is +/- 1.2 percentage points.

Table 8: Bootstrapped Confidence Intervals Computed on Assumption of Normality

Variable	Mean	Std. Dev.	Confidence interval	
			LB	UB
Total Accuracy	85.5969	0.5614	84.4966	86.6973
Poverty Accuracy	59.9915	1.7434	56.5743	63.4088
Undercoverage	40.0084	1.7434	36.5912	43.4257
Leakage	36.0337	2.4422	31.2468	40.8206
PIE	-0.7617	0.5872	-1.9127	0.3891
BPAC	55.7572	3.9546	48.0062	63.5082

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 2.5th percentile of the sample distribution for each measure; the upper bound is defined by the 97.5th 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	95% Confidence Interval	
	LB	UB
Total Accuracy	84.5007	86.6792
Poverty Accuracy	56.4477	63.3470
Undercoverage	36.6530	43.5523
Leakage	31.4665	41.1394
PIE	-2.0084	0.3870
BPAC	46.7725	62.1118

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 Ghana PAT appears to be somewhat biased toward underestimating the actual extreme poverty rate, but nonetheless will fall within 2 percentage points of the true value in the population (with 95 percent confidence). By this measure, the predictive model behind the Ghana PAT is accurate.