Poverty Assessment Tool Accuracy Submission
USAID/IRIS Tool for Kenya
Submitted: July 20, 2010
Out-of-sample bootstrap results added: October 20, 2010
Typo corrected: July 31, 2012

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

1. Data source

For Kenya, existing data from the 2005/2006 Integrated Household Budget Survey (KIHBS) were used to construct the poverty assessment tool. The full sample of 13,158 households is nationally representative. The sample used for tool construction comprises a randomly selected 9,868 households (75 percent of the full sample). The remainder, another randomly selected 3,290 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, illness and disability, and employment. For Kenya, more than 100 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 Kenya 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 Kenya, the most accurate method, on the basis of BPAC, is the 2-step LP 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

KENYA $1.25/day line***** Share of "very poor": 36.3%	**Total Accuracy**	**Poverty Accuracy**	**Under-coverage**	**Leakage**	**PIE**	**BPAC**
Single-step methods --						
OLS	79.31	70.57	29.43	27.97	-0.53	69.12
Quantile regression (estimation point: 49 percentile)	**79.41**	**71.35**	**28.65**	**28.49**	**-0.06**	**71.19**
Linear Probability	79.73	70.36	29.64	26.59	-1.10	67.31
Probit	79.80	69.84	30.16	25.89	-1.54	65.57
Two-step methods --						
OLS – 67 percentile cutoff	80.26	72.85	27.15	27.61	0.17	72.39
Quantile (estimation points: 49, 33) 67 percentile cutoff	80.06	72.67	27.33	27.98	0.23	72.02
LP – 57 percentile cutoff	80.47	72.77	27.23	26.95	-0.10	72.48
Probit –57 percentile cutoff	80.25	71.68	28.32	26.47	-0.66	69.84

For Kenya, 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.

[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.

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 Kenya 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. The value of the line in those prices is 1,323 Kenyan Shillings per capita per month.[4] At these values, the $1.25/day poverty line identifies 36.3% of households as "very poor." This compares with an estimate from PovcalNet of 19.7%. Substantial effort was made to resolve this difference. According to direct communication with the Kenyan NSO, IRIS is computing this measure using the correct household expenditure series and household size to derive the appropriate per capita consumption figure.

Alternatively, the national poverty line of 1,562 Kenyan Shillings per adult equivalent per month in rural areas and 2,913 in urban areas identifies 38.3% of households as poor and therefore 19.1% as "very poor." This matches external sources.

Hence the decision rule for Kenya'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 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 Kenya is 32.68.

[4] The calculation for the $1.25/day poverty line is (32.6837465718227*1.25)*(365.25/12)*(106.4/100) 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)	2,538 (25.7%)	1,013 (10.3%)
Number of "true" not very-poor households (as determined by benchmark survey)	1,019 (10.3%)	5,298 (53.7%)

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

.49 Quantile regression Number of obs = 9868
Min sum of deviations 4212.144 Pseudo R2 = 0.3746

Variable	Coef.	Std. Err.	t	P>|t|	[95% Conf. Interval]	
Intercept	8.1973	0.0795	103.0700	0.0000	8.0414	8.3532
Household size	-0.3132	0.0033	2.1000	0.0360	0.0005	0.0136
Household size squared	0.0146	0.0000	-3.0500	0.0020	-0.0002	0.0000
Household head age	0.0070	0.0095	-33.0600	0.0000	-0.3318	-0.2946
Household head age squared	-0.0001	0.0007	21.7200	0.0000	0.0133	0.0159
HH lives in rural area	-0.3211	0.0279	-11.5100	0.0000	-0.3758	-0.2664
HH lives in central region	0.0152	0.0291	0.5200	0.6020	-0.0418	0.0721
HH lives in coast region	0.0054	0.0330	0.1600	0.8690	-0.0592	0.0701
HH lives in eastern region	-0.0199	0.0265	-0.7500	0.4510	-0.0718	0.0319
HH lives in Nairobi	0.1306	0.0458	2.8500	0.0040	0.0409	0.2203
HH lives in nyanza region	-0.1124	0.0270	-4.1600	0.0000	-0.1654	-0.0595
HH lives in western region	-0.1145	0.0306	-3.7400	0.0000	-0.1746	-0.0545
HH lives in north eastern region	0.0210	0.0440	0.4800	0.6340	-0.0652	0.1072
Household head is female	-0.0841	0.0186	-4.5300	0.0000	-0.1204	-0.0477
HH main source of cooking fuel is purchased firewood	0.1411	0.0267	5.2900	0.0000	0.0888	0.1935
HH main source of drinking water over the past month is water piped into dwelling	0.2709	0.0432	6.2700	0.0000	0.1862	0.3556
HH main source of drinking water over the past month is water piped into plot or yard	0.1288	0.0302	4.2600	0.0000	0.0695	0.1881
Dwelling main flooring material is cement	0.2348	0.0223	10.5200	0.0000	0.1911	0.2785
HH member raised or owned livestock, poultry, fish, etc. in last 12 months	0.0632	0.0224	2.8200	0.0050	0.0192	0.1071
Number of refrigerators owned by HH	0.3504	0.0538	6.5100	0.0000	0.2449	0.4560
Number of electric irons owned by HH	0.1687	0.0361	4.6700	0.0000	0.0978	0.2395
Number of jiko-charcol owned by HH	0.0890	0.0139	6.3900	0.0000	0.0617	0.1163
Number of radios owned by HH	0.0848	0.0130	6.5400	0.0000	0.0594	0.1102
Number of rooms occupied by HH	0.0812	0.0077	10.5200	0.0000	0.0661	0.0963
HH owns one or more electric or gas cookers	0.3118	0.0491	6.3500	0.0000	0.2156	0.4081
HH owns one or more charcoal irons	0.1890	0.0226	8.3500	0.0000	0.1446	0.2333
HH owns one or more radios with cassette or CD player	0.1626	0.0226	7.1800	0.0000	0.1182	0.2070
HH toilet facility is a bucket or none	-0.0881	0.0235	-3.7400	0.0000	-0.1342	-0.0420

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 Kenya are defined as those whose predicted expenditures fall below (above) the $1.25/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 Kenya, 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 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

Kenya $2.50/day Line Share of Poor: 68.3%	**Total Accuracy**	**Poverty Accuracy**	**Under-coverage**	**Leakage**	**PIE**	**BPAC**
Single-step methods						
Quantile regression (estimation point: 56)	83.97	88.50	11.50	11.96	0.31	88.04

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 poor by the tool	**Number of households identified as not poor by the tool**
Number of "true" poor households (as determined by benchmark survey)	5,967 (60.4%)	806 (8.2%)
Number of "true" not poor households (as determined by benchmark survey)	775 (7.9%)	2,319 (23.5%)

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

.56 Quantile regression — Number of obs = 9868

Min sum of deviations 4153.817 — Pseudo R2 = 0.3829

Variable	Coef.	Std. Err.	t	P>|t|	[95% Conf. Interval]	
Intercept	8.3296	0.0749	111.1800	0.0000	8.1828	8.4765
Household size	-0.3170	0.0031	2.2800	0.0230	0.0010	0.0132
Household size squared	0.0149	0.0000	-3.3700	0.0010	-0.0002	0.0000
Household head age	0.0071	0.0088	-36.1900	0.0000	-0.3342	-0.2999
Household head age squared	-0.0001	0.0006	24.6600	0.0000	0.0137	0.0161
HH lives in rural area	-0.3462	0.0261	-13.2800	0.0000	-0.3974	-0.2951
HH lives in central region	0.0302	0.0274	1.1100	0.2690	-0.0234	0.0839
HH lives in coast region	-0.0145	0.0310	-0.4700	0.6400	-0.0753	0.0463
HH lives in eastern region	-0.0107	0.0248	-0.4300	0.6650	-0.0594	0.0379
HH lives in Nairobi	0.1258	0.0427	2.9500	0.0030	0.0421	0.2094
HH lives in nyanza region	-0.0915	0.0253	-3.6200	0.0000	-0.1411	-0.0419
HH lives in western region	-0.1201	0.0288	-4.1700	0.0000	-0.1766	-0.0636
HH lives in north eastern region	0.0333	0.0416	0.8000	0.4240	-0.0484	0.1149
Household head is female	-0.0850	0.0174	-4.8900	0.0000	-0.1190	-0.0509
HH main source of cooking fuel is purchased firewood	0.1327	0.0251	5.2900	0.0000	0.0835	0.1819
HH main source of drinking water over the past month is water piped into dwelling	0.2158	0.0406	5.3100	0.0000	0.1362	0.2954
HH main source of drinking water over the past month is water piped into plot or yard	0.1086	0.0283	3.8400	0.0000	0.0532	0.1640
Dwelling main flooring material is cement	0.2188	0.0211	10.3700	0.0000	0.1774	0.2602
HH member raised or owned livestock, poultry, fish, etc. in last 12 months	0.0556	0.0212	2.6300	0.0090	0.0141	0.0971
Number of refrigerators owned by HH	0.4869	0.0497	9.8000	0.0000	0.3895	0.5843
Number of electric irons owned by HH	0.1431	0.0333	4.3000	0.0000	0.0778	0.2083
Number of jiko-charcol owned by HH	0.0822	0.0132	6.2400	0.0000	0.0564	0.1080
Number of radios owned by HH	0.0865	0.0122	7.0700	0.0000	0.0625	0.1104
Number of rooms occupied by HH	0.0824	0.0071	11.5600	0.0000	0.0684	0.0963
HH owns one or more electric or gas cookers	0.3291	0.0452	7.2900	0.0000	0.2406	0.4176
HH owns one or more charcoal irons	0.1770	0.0212	8.3600	0.0000	0.1355	0.2186
HH owns one or more radios with cassette or CD player	0.1695	0.0212	8.0100	0.0000	0.1280	0.2110
HH toilet facility is a bucket or none	-0.0888	0.0221	-4.0200	0.0000	-0.1320	-0.0455

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 Kenya tool was divided randomly into two data sets 9,868 households (75 percent of the sample) and 3,290 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 slightly more than 2 BPAC points and 0.8 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						
	79.41	71.35	28.65	28.49	-0.06	71.19
Out-of-Sample Prediction						
	79.69	71.43	28.57	26.16	-0.89	69.02

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 +/- 4.8 percentage points. For PIE,

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

which measures the difference between the predicted poverty rate and the actual poverty rate, the confidence interval is +/- 2.1 percentage points.

Table 8: Bootstrapped Confidence Intervals on Assumption of Normality

Variable	Mean	Std. Dev.	Confidence interval	
			LB	UB
Total Accuracy	79.32	0.91	77.54	81.10
Poverty Accuracy	72.08	1.68	68.79	75.36
Undercoverage	27.93	1.68	24.64	31.21
Leakage	27.84	2.14	23.65	32.02
PIE	-0.04	1.03	-2.07	1.98
BPAC	69.88	2.47	65.04	74.72

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 indicate that the extreme poverty rate estimate produced by the Kenya PAT is unbiased and will fall within 2.1 percentage points of the true value in the population (with 95 percent confidence). By this measure, the predictive model behind the Kenya PAT has a high degree of accuracy.