Poverty Assessment Tool Accuracy Submission
USAID/IRIS Tool for Rwanda
Submitted: June 30, 2011

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

1. Data source

For Rwanda, existing data from the 2005/2006 Enquete Intégrale sur les Conditions de Vie des ménages de Rwanda (EICV2) were used to construct the poverty assessment tool. The full sample of 6,894 households is nationally representative. The sample used for tool construction comprises a randomly selected 5,171 households (75 percent of the full sample). The remainder, another randomly selected 1,723 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. The data collected in the household roster was used to construct gender indicators such as *household head is female*, *share of female household members of age 6-15*, and *share of female household members of age 16-25*. For Rwanda, more than 82 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 Rwanda 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 Rwanda, on the basis of BPAC, the 1-step and 2-step Quantile regression models are equally accurate methods. However, the 1-step Quantile regression 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

Rwanda ($1.25/day line*) Share of "very poor": 74.7%	**Total Accuracy**	**Poverty Accuracy**	**Under-coverage**	**Leakage**	**PIE**	**BPAC**
Single-step methods						
OLS	83.40	93.57	6.43	16.37	7.23	83.63
Quantile regression (estimation point: 64 percentile)	**81.98**	**87.69**	**12.31**	**12.45**	**0.11**	**87.55**
Linear Probability	83.29	95.92	4.08	18.89	10.78	81.11
Probit	83.32	94.52	5.48	17.44	8.70	82.56
Two-step methods						
OLS –73 percentile cutoff	83.19	92.11	7.89	15.20	5.32	84.80
Quantile (estimation points: 64 , 47) 73 percentile cutoff	81.98	87.69	12.31	12.45	0.11	87.55
LP – 73 percentile cutoff	83.46	95.76	4.24	18.48	10.36	81.52
Probit –73 percentile cutoff	83.36	93.82	6.18	16.68	7.64	83.32
*The $1.25 per day per capita international poverty line in 2005 Purchasing Power Parity terms is 322 Rwandan Francs per day per capita in 2006 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 Rwanda, 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 Rwanda 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 (2005-2006). The value of the line in those prices is 322 Rwandan Francs per day per capita.[4] At these values, the \$1.25/day poverty line identifies 72.9% of households as "very poor."

By comparison, in 2005-2006 53.9% of Rwandan households lived below the national poverty line of 90,000 Rwandan Francs per adult equivalent per year. According to the second definition of extreme poverty in the legislation, the median poverty line of 55,635 Rwandan Francs per adult equivalent per year identifies 26.9% of households as "very poor."

Hence the decision rule for Rwanda'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 Rwanda is 236.75.

[4] The calculation for the \$1.25/day poverty line is 1.25*(236.74524)*(108.883/100.0) where the final term is the CPI adjustment from average 2005 prices to average 2006 prices.

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)	3,301 (63.8%)	463 (9.0%)
Number of "true" not very-poor households (as determined by benchmark survey)	469 (9.1%)	938 (18.1%)

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

.64 Quantile regression — Number of obs = 5,171
Min sum of deviations 2087.207 — Pseudo R2 = 0.3403

Variable	Coef.	Std. Err.	t	P>|t|	[95% Conf. Interval]	
Intercept	6.2898	0.1078	58.3500	0.0000	6.0785	6.5011
Household size	-0.2461	0.0171	-14.3900	0.0000	-0.2796	-0.2126
Household size squared	0.0119	0.0014	8.7700	0.0000	0.0092	0.0145
Household head age	-0.0014	0.0040	-0.3400	0.7340	-0.0093	0.0066
Household head age squared	0.0000	0.0000	0.0200	0.9840	-0.0001	0.0001
Household lives in rural area	-0.1436	0.0359	-4.0000	0.0000	-0.2140	-0.0732
Household lives in Kigali-Ngali	-0.2155	0.0581	-3.7100	0.0000	-0.3294	-0.1016
Household lives in Gitarama	-0.3469	0.0565	-6.1400	0.0000	-0.4577	-0.2361
Household lives in Butare	-0.4767	0.0581	-8.2000	0.0000	-0.5906	-0.3628
Household lives in Gikongoro	-0.4072	0.0617	-6.6000	0.0000	-0.5282	-0.2862
Household lives in Cyangugu	-0.2775	0.0594	-4.6700	0.0000	-0.3939	-0.1611
Household lives in Kibuye	-0.2387	0.0591	-4.0400	0.0000	-0.3546	-0.1228
Household lives in Gisenyi	-0.1838	0.0569	-3.2300	0.0010	-0.2954	-0.0722
Household lives in Ruhengeri	-0.2809	0.0579	-4.8500	0.0000	-0.3943	-0.1674
Household lives in Byumba	-0.3237	0.0590	-5.4900	0.0000	-0.4393	-0.2081
Household lives in Umutara	-0.0392	0.0617	-0.6400	0.5250	-0.1601	0.0817
Household lives in Kibungo	-0.1273	0.0587	-2.1700	0.0300	-0.2425	-0.0121
Roof of dwelling is made of thatch	-0.1255	0.0366	-3.4300	0.0010	-0.1972	-0.0538
Floor of dwelling is made of cement	0.3335	0.0392	8.5100	0.0000	0.2567	0.4104
HH electricity source is public utility	0.8036	0.0576	13.9400	0.0000	0.6906	0.9167
HH electricity source is kerosene lantern	0.1860	0.0345	5.3800	0.0000	0.1183	0.2537
HH electricity source is firewood	-0.2929	0.0316	-9.2800	0.0000	-0.3548	-0.2310
Toilet is an unprotected latrine	-0.0755	0.0225	-3.3500	0.0010	-0.1197	-0.0313
Main cooking fuel is charcoal	0.4102	0.0520	7.8800	0.0000	0.3082	0.5122
Number of rooms in dwelling	0.0703	0.0088	8.0400	0.0000	0.0532	0.0875
Share of members age 0-5	-0.4596	0.0672	-6.8400	0.0000	-0.5914	-0.3278
Household head has no schooling	-0.0918	0.0239	-3.8400	0.0000	-0.1387	-0.0449
Household head is currently employed	0.1366	0.0377	3.6200	0.0000	0.0626	0.2106
Number of radio or cassette players owned	0.1439	0.0178	8.1000	0.0000	0.1090	0.1787
Number of bicycles owned	0.2055	0.0299	6.8700	0.0000	0.1468	0.2641
Number of hoe or shovels owned	0.0508	0.0088	5.7900	0.0000	0.0336	0.0681
Household owns one or more heads of cattle	0.1487	0.0264	5.6300	0.0000	0.0970	0.2005

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 Rwanda 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 Rwanda, 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

Rwanda ($2.50/day line*) Share of "very poor": 90.8%	**Total Accuracy**	**Poverty Accuracy**	**Under-coverage**	**Leakage**	**PIE**	**BPAC**
Single-step method						
Quantile regression (estimation point: 70)	93.11	96.22	3.78	3.81	0.03	96.19
*The $2.50 per day per capita poverty line in 2005 Purchasing Power Parity terms is 644 Rwandan Francs per day per capita in 2006 prices.						

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)	4,514 (87.3%)	177 (3.4%)
Number of "true" not poor households (as determined by benchmark survey)	179 (3.5%)	301 (5.8%)

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

.70 Quantile regression — Number of obs = 5,171
Min sum of deviations 1961.189 — Pseudo R2 = 0.3480

Variable	Coef.	Std. Err.	t	P>|t|	[95% Conf. Interval]	
Intercept	6.40158	0.08916	71.8000	0.0000	6.2268	6.5764
Household size	-0.24821	0.01382	-17.9600	0.0000	-0.2753	-0.2211
Household size squared	0.01199	0.00109	11.0300	0.0000	0.0099	0.0141
Household head age	-0.00132	0.00333	-0.4000	0.6920	-0.0079	0.0052
Household head age squared	0.00001	0.00003	0.1900	0.8490	-0.0001	0.0001
Household lives in rural area	-0.13217	0.02947	-4.4800	0.0000	-0.1899	-0.0744
Household lives in Kigali-Ngali	-0.24672	0.04730	-5.2200	0.0000	-0.3394	-0.1540
Household lives in Gitarama	-0.36842	0.04630	-7.9600	0.0000	-0.4592	-0.2776
Household lives in Butare	-0.47549	0.04764	-9.9800	0.0000	-0.5689	-0.3821
Household lives in Gikongoro	-0.44864	0.05000	-8.9700	0.0000	-0.5467	-0.3506
Household lives in Cyangugu	-0.28382	0.04867	-5.8300	0.0000	-0.3792	-0.1884
Household lives in Kibuye	-0.27467	0.04865	-5.6500	0.0000	-0.3700	-0.1793
Household lives in Gisenyi	-0.20732	0.04665	-4.4400	0.0000	-0.2988	-0.1159
Household lives in Ruhengeri	-0.30803	0.04735	-6.5100	0.0000	-0.4009	-0.2152
Household lives in Byumba	-0.33884	0.04826	-7.0200	0.0000	-0.4335	-0.2442
Household lives in Umutara	-0.06606	0.05068	-1.3000	0.1920	-0.1654	0.0333
Household lives in Kibungo	-0.13401	0.04787	-2.8000	0.0050	-0.2279	-0.0402
Roof of dwelling is made of thatch	-0.13465	0.02970	-4.5300	0.0000	-0.1929	-0.0764
Floor of dwelling is made of cement	0.33675	0.03231	10.4200	0.0000	0.2734	0.4001
HH electricity source is public utility	0.88337	0.04704	18.7800	0.0000	0.7911	0.9756
HH electricity source is kerosene lantern	0.19225	0.02799	6.8700	0.0000	0.1374	0.2471
HH electricity source is firewood	-0.31987	0.02569	-12.4500	0.0000	-0.3702	-0.2695
Toilet is an unprotected latrine	-0.06504	0.01858	-3.5000	0.0000	-0.1015	-0.0286
Main cooking fuel is charcoal	0.36802	0.04269	8.6200	0.0000	0.2843	0.4517
Number of rooms in dwelling	0.06574	0.00711	9.2400	0.0000	0.0518	0.0797
Share of members age 0-5	-0.46544	0.05554	-8.3800	0.0000	-0.5743	-0.3566
Household head has no schooling	-0.08512	0.01965	-4.3300	0.0000	-0.1236	-0.0466
Household head is currently employed	0.14093	0.03119	4.5200	0.0000	0.0798	0.2021
Number of radio or cassette players owned	0.13836	0.01474	9.3900	0.0000	0.1095	0.1672
Number of bicycles owned	0.20387	0.02450	8.3200	0.0000	0.1558	0.2519
Number of hoe or shovels owned	0.04889	0.00719	6.8000	0.0000	0.0348	0.0630
Household owns one or more heads of cattle	0.12701	0.02173	5.8400	0.0000	0.0844	0.1696

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 Rwanda tool was divided randomly into two data sets 5,171 households (75 percent of the sample) and 1,723 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, losing 0.38 BPAC points and 0.02 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						
	81.98	87.69	12.31	12.45	0.11	87.55
Out-of-Sample Prediction						
	81.59	87.35	12.65	12.46	-0.13	87.17

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 extremely narrow at +/- 2.8 percentage points. For PIE, 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	81.58	0.97	79.67	83.49
Poverty Accuracy	87.34	0.98	85.43	89.26
Undercoverage	12.66	0.98	10.74	14.57
Leakage	12.47	1.08	10.36	14.58
PIE	-0.14	1.06	-2.22	1.93
BPAC	86.18	1.39	83.46	88.90

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	79.66	83.50
Poverty Accuracy	85.43	89.24
Undercoverage	10.76	14.57
Leakage	10.53	14.76
PIE	-2.17	2.00
BPAC	82.86	88.30

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 Rwanda PAT has a confidence interval of +/- 2.1, which is relatively narrow. By this measure, the predictive model behind the Rwanda PAT is accurate