Poverty Assessment Tool Accuracy Submission
USAID/IRIS Tool for Ecuador
Submitted: June 30, 2010

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

1. Data source

For Ecuador, existing data from the 2005/2006 Encuesta de Condiciones de Vida (ECV) were used to construct the poverty assessment tool. The full sample of 13,581 households was not used for tool construction because household daily per capita consumption was missing, without a discernible pattern, for 149 households. The remaining sample of 13,432 households is nationally representative. The sample used for tool construction comprises a randomly selected 10,074 households (75 percent of the sample). The remainder, another randomly selected 3,358 households (25 percent of the sample), is reserved for out-of-sample accuracy testing, which investigates 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 Ecuador, nearly 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 Ecuador is the median 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 Ecuador, 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

ECUADOR Median line***** Share of "very poor": 14.8%	**Total Accuracy**	**Poverty Accuracy**	**Under-coverage**	**Leakage**	**PIE**	**BPAC**
Single-step methods --						
OLS	89.35	53.12	46.88	25.26	-3.19	31.50
Quantile regression (estimation point: 39 percentile)	**88.55**	**61.89**	**38.11**	**39.47**	**0.20**	**60.53**
Linear Probability	89.33	38.95	61.05	11.22	-7.35	-10.88
Probit	89.99	49.81	50.19	17.60	-4.81	17.23
Two-step methods --						
OLS –99 percentile cutoff	89.33	53.13	46.87	25.45	-3.16	31.71
Quantile (estimation points: 39, 40) 99 percentile cutoff	88.69	61.47	38.53	38.12	-0.06	61.05
LP – 24 percentile cutoff	90.26	53.48	46.52	19.47	-3.99	26.42
Probit –24 percentile cutoff	89.97	52.65	47.35	20.58	-3.95	25.87
* Median poverty line is US$ 1.32 per capita per day in 2006 prices. This median poverty line is based on the official national poverty line of US$ 1.89 per capita per day.						

[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 Ecuador, the functionality of predicting the poverty rate at other poverty lines—in this case, the food poverty line, halfway between the median and national poverty line, the national line, and 50% above the national line —has been added. This functionality is based on statistical models for prediction at the median and national 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 Ecuador the applicable threshold is the median poverty line of US$1.32 per capita per day, the household per capita expenditure value of the 50th percentile below the national poverty line,[4] at the level of prices prevailing when the household survey data were collected. At these values, the median poverty line identifies 14.8% of households as "very poor".

Alternatively, the international poverty line of $1.25/day in 2005 PPP terms line identifies 2.1% of households as "very poor."[5]

The national poverty line of $US 1.89 per capita per day yields a household poverty rate of 29.7%.

---

[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 Ecuadoris 4475.822.

[4] The currency in Ecuador is the U.S. dollar.

[5] The World Bank's Povcalnet provides a poverty headcount of 9.7%. However, this number is based on a labor and employment survey and considers income not consumption expenditures. In any case, the $1.25/day PPP line would not be binding.

Hence the decision rule for Ecuador'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 median 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 median poverty line is identified as "very poor," and each household whose estimated per capita consumption expenditures exceeds the median 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)	920 (9.13%)	567 (5.63%)
Number of "true" not very-poor households (as determined by benchmark survey)	587 (5.83%)	8000 (79.41%)

Table 3: Regression Estimates using 1-step Quantile Method for Prediction at the Median Poverty Line

ECUADOR 1-STEP MAXR/QUANT: variables from MAXR/OLS 100 percentile model
Regression results, estimation point of 39 percentile

.39 Quantile regression — Number of obs = 10074
Min sum of deviations 2988.481 — Pseudo R2 = 0.5186

Variable	Coef.	Std. Err.	t	P>|t|	[95% Conf. Interval]	
Intercept	0.5813	0.0481	12.0900	0.0000	0.4871	0.6756
Household size	-0.3245	0.0074	-44.1500	0.0000	-0.3389	-0.3101
Household size squared	0.0139	0.0006	21.5800	0.0000	0.0126	0.0151
Household head age	0.0118	0.0018	6.4800	0.0000	0.0082	0.0153
Household head age squared	-0.0001	0.0000	-6.7700	0.0000	-0.0002	-0.0001
Household lives in Costa region	0.0444	0.0116	3.8300	0.0000	0.0217	0.0670
Household lives in Amazonia region	-0.0434	0.0190	-2.2800	0.0230	-0.0808	-0.0061
Household lives in rural area	-0.0896	0.0124	-7.2100	0.0000	-0.1139	-0.0652
Final degree obtained by household head is a high school diploma	0.1036	0.0151	6.8800	0.0000	0.0741	0.1331
Final degree obtained by household head is a university degree	0.2574	0.0212	12.1600	0.0000	0.2159	0.2989
Share of household members with no schooling	-0.2319	0.0269	-8.6300	0.0000	-0.2846	-0.1792
Floor of dwelling is finished wood or parquet	0.2235	0.0185	12.1000	0.0000	0.1873	0.2597
Floor of dwelling is ceramic or vinyl tile	0.2032	0.0162	12.5600	0.0000	0.1715	0.2349
Number of rooms in dwelling	0.0788	0.0043	18.3300	0.0000	0.0704	0.0872
Household owns one or stoves	0.3306	0.0228	14.4800	0.0000	0.2859	0.3754
Household owns one or more refrigerators	0.1538	0.0138	11.1400	0.0000	0.1268	0.1809
Number of blenders HH owns	0.1363	0.0129	10.5300	0.0000	0.1109	0.1617
Household owns one or more irons	0.1396	0.0143	9.7600	0.0000	0.1115	0.1676
Number of color TVs HH owns	0.0919	0.0081	11.4000	0.0000	0.0761	0.1077
Number of cars HH owns	0.3648	0.0131	27.8400	0.0000	0.3391	0.3905
Number of computers HH owns	0.1383	0.0158	8.7500	0.0000	0.1073	0.1693
Number of washing machines HH owns	0.1325	0.0159	8.3300	0.0000	0.1013	0.1637
Household owns one or more hand mixers	0.0873	0.0144	6.0400	0.0000	0.0590	0.1156

Annex 1: Poverty Prediction at the National 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 Ecuador are defined as those whose predicted incomes fall below (above) the national in 2005 PPP.

Table 4 summarizes the predictive accuracy results for the national poverty line using the Quantile model specification from the median poverty line. The indicators are the same as those in the model for the median 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 national 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 Ecuador, these five lines are the food poverty line, the median line, halfway between the median and national poverty line, the national line, and 50% above the national line. Poverty rates at the first three lines are predicted using the best model for the median line, while poverty rates at the last two lines are predicted using the best model for the national line. As discussed in this document, accuracy has been tested at the median and national 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 National Poverty Line

Ecuador National Line Share of Poor: 29.7%	**Total Accuracy**	**Poverty Accuracy**	**Under-coverage**	**Leakage**	**PIE**	**BPAC**
Single-step methods						
Quantile regression (estimation point: 43)	84.75	74.40	25.60	25.33	-0.08	74.13

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 National 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)	2244 (22.3%)	772 (7.7%)
Number of "true" not poor households (as determined by benchmark survey)	764 (7.6%)	6294 (62.5%)

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

ECUADOR 1-STEP MAXR/QUANT: variables from MAXR/OLS 100 percentile model
Regression results, estimation point of 43 percentile

.43 Quantile regression — Number of obs = 10074
Min sum of deviations 3056.503 — Pseudo R2 = 0.5225

Variable	**Coef.**	**Std. Err.**	**t**	**P>|t|**	**[95% Conf. Interval]**	
Intercept	0.6401	0.0558	11.4700	0.0000	0.5307	0.74950
Household size	-0.3319	0.0083	-39.9700	0.0000	-0.3482	-0.3156
Household size squared	0.0145	0.0007	20.3400	0.0000	0.0131	0.01587
Household head age	0.0119	0.0021	5.6600	0.0000	0.0078	0.01607
Household head age squared	-0.0001	0.0000	-5.8900	0.0000	-0.0002	-0.0001
Household lives in Costa region	0.0384	0.0134	2.8700	0.0040	0.0122	0.0647
Household lives in Amazonia region	-0.0362	0.0220	-1.6400	0.1000	-0.0794	0.0070
Household lives in rural area	-0.0779	0.0144	-5.4200	0.0000	-0.1060	-0.0497
Final degree obtained by household head is a high school diploma	0.1135	0.0174	6.5100	0.0000	0.0794	0.1477
Final degree obtained by household head is a university degree	0.2634	0.0245	10.7300	0.0000	0.2153	0.3115
Share of household members with no schooling	-0.2369	0.0311	-7.6200	0.0000	-0.2978	-0.1760
Floor of dwelling is finished wood or parquet	0.2210	0.0214	10.3300	0.0000	0.1791	0.2629
Floor of dwelling is ceramic or vinyl tile	0.1901	0.0188	10.1400	0.0000	0.1533	0.2269
Number of rooms in dwelling	0.0769	0.0050	15.5200	0.0000	0.0672	0.0866
Household owns one or stoves	0.3380	0.0265	12.7700	0.0000	0.2861	0.3899
Household owns one or more refrigerators	0.1514	0.0160	9.4700	0.0000	0.1200	0.1827
Number of blenders HH owns	0.1286	0.0150	8.6000	0.0000	0.0993	0.1579
Household owns one or more irons	0.1434	0.0165	8.6600	0.0000	0.1109	0.1758
Number of color TVs HH owns	0.0879	0.0093	9.4200	0.0000	0.0696	0.1062
Number of cars HH owns	0.3690	0.0151	24.4600	0.0000	0.3394	0.3985
Number of computers HH owns	0.1454	0.0183	7.9500	0.0000	0.1095	0.1812
Number of washing machines HH owns	0.1298	0.0185	7.0300	0.0000	0.0936	0.1660
Household owns one or more hand mixers	0.0963	0.0168	5.7400	0.0000	0.0634	0.1291

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 Ecuador tool was divided randomly into two data sets of 10,074 households (75 percent of the sample) and 3,358 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, actually gaining 1.7 points in BPAC and losing 0.08 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						
	88.55	61.89	38.11	39.47	0.20	60.53
Out-of-Sample Prediction						
	89.33	64.05	35.95	34.12	-0.28	62.22

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.[6] 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.

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

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 +/- 8.4 percentage points. For PIE, which measures the difference between the predicted poverty rate and the actual poverty rate, the confidence interval is +/- 1.3 percentage points.

Table 8: Bootstrapped Confidence Intervals on Assumption of Normality

Variable	Mean	Std. Dev.	Confidence interval	
			LB	UB
Total Accuracy	89.3374	0.6215	88.1193	90.5555
Poverty Accuracy	64.0651	2.4638	59.2360	68.8942
Undercoverage	35.9349	2.4638	31.1058	40.7640
Leakage	34.0486	3.4136	27.3580	40.7392
PIE	-0.3024	0.6477	-1.5719	0.9671
BPAC	60.3604	4.2594	52.0119	68.7088

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	88.1484	90.4951
Poverty Accuracy	59.0698	68.5571
Undercoverage	31.4429	40.9302
Leakage	27.8242	40.7976
PIE	-1.6301	0.9531
BPAC	49.9611	67.0378

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 Ecuador PAT appears to be slightly 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 Ecuador PAT is accurate.