

**Poverty Assessment Tool Accuracy Submission**  
**USAID/IRIS Tool for El Salvador**  
**Submitted: July 18, 2010**

The following report is divided into five sections. Section 1 describes the data set used to create the Poverty Assessment Tool for El Salvador. Section 2 details the set of statistical procedures used for selecting indicators and for estimating household income 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 El Salvador Poverty Assessment Tool.

**1. Data source**

For El Salvador, existing data from the 2008 Encuesta de Hogares de Propósitos Múltiples (EHPM) were used to construct the poverty assessment tool. The full sample of 16,674 households is nationally representative. The sample used for tool construction comprises a randomly selected 12,506 households (75 percent of the full sample). The remainder, another randomly selected 4,168 households, 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 El Salvador, nearly 80 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.<sup>1</sup> 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.<sup>2</sup>

### 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 El Salvador 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 El Salvador, 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**

<b>El Salvador</b> Median line* Share of “very poor”: 20.2%	<b>Total Accuracy</b>	<b>Poverty Accuracy</b>	<b>Under-coverage</b>	<b>Leakage</b>	<b>PIE</b>	<b>BPAC</b>
<b>Single-step methods</b>						
OLS	82.30	36.49	63.51	26.12	-7.38	-0.89
<b>Quantile regression (estimation point: 38 percentile)</b>	<b>81.10</b>	<b>52.84</b>	<b>47.16</b>	<b>46.83</b>	<b>-0.07</b>	<b>52.51</b>
Linear Probability	82.46	25.40	74.60	14.22	-11.92	-34.97
Probit	82.59	33.37	66.63	21.52	-8.91	-11.74
<b>Two-step methods</b>						
OLS –39 percentile cutoff	82.96	38.50	61.50	24.78	-7.25	1.78
Quantile (estimation points: 38, 15) 39 percentile cutoff	81.70	54.11	45.90	45.08	-0.16	53.29
LP – 41 percentile cutoff	82.77	34.12	65.88	21.39	-8.79	-10.37
Probit – 41 percentile cutoff	82.91	33.95	66.05	20.46	-9.00	-11.63
*Median poverty line is US\$ 1.97 per capita per day in 2008 prices in urban areas and US\$ 1.17 in rural areas. This poverty line is based on the official national poverty line of US\$ 2.94 per capita per day in urban areas and US\$ 1.91 in rural areas						

For El Salvador, 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 —have been added. This functionality

<sup>1</sup> The set of indicators and their weights also depended on the selection of a 1-step or 2-step statistical model.

<sup>2</sup> For a detailed discussion of these accuracy criteria, see “Note on Assessment and Improvement of Tool Accuracy” at [www.povertytools.org](http://www.povertytools.org).

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 income**

For the quantile regression method, the estimated regression coefficients indicate the weight placed on each of the included indicators in estimating the household income 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)<sup>3</sup>. The applicable poverty line for USAID tool development is the one that yields the higher household poverty rate for a given country.

In El Salvador the applicable threshold is the median poverty line, the household per capita income value of the 50<sup>th</sup> percentile below the national poverty line, of US\$ 1.97 per capita per day in urban areas and US\$ 1.17 in rural areas, at the level of prices prevailing in 2008 when the household survey data were collected. At these values, the median poverty line identifies 20.2% of households as “very poor.”

Alternatively, the international poverty line of \$1.25/day in 2005 PPP terms identifies 5.4% of households as “very poor.”<sup>4</sup>

Hence the decision rule for El Salvador’s USAID poverty assessment tool in classifying the “very poor” (and the “not very-poor”) is whether that predicted per capita daily income 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 income according to the tool is less than or equal to the median

---

<sup>3</sup> 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 El Salvador is 4.812.

<sup>4</sup> The World Bank’s Povcalnet provides a poverty headcount of 11.0% using population weights and a different data set.

poverty line is identified as “very poor,” and each household whose estimated per capita consumption income 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**

	<b>Number of households identified as very poor by the tool</b>	<b>Number of households identified as not very-poor by the tool</b>
<b>Number of “true” very poor households (as determined by benchmark survey)</b>	1329 (10.6%)	1186 (9.5%)
<b>Number of “true” not very-poor households (as determined by benchmark survey)</b>	1178 (9.4)	8813 (70.5 %)

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

38 Quantile regression  
 Min sum of deviations 5328.474

Number of obs = 12506  
 Pseudo R2 = 0.3430

<b>Variable</b>	<b>Coef.</b>	<b>Std. Err.</b>	<b>t</b>	<b>P&gt; t </b>	<b>[95% Conf. Interval]</b>	
Intercept	0.6372	0.0592	10.7700	0.0000	0.5213	0.7532
Household size	-0.3108	0.0097	-32.2100	0.0000	-0.3297	-0.2919
Household size squared	0.0176	0.0009	20.4200	0.0000	0.0159	0.0193
Household head age	0.0162	0.0023	7.1500	0.0000	0.0118	0.0207
Household head age squared	-0.0001	0.0000	-6.6500	0.0000	-0.0002	-0.0001
Household lives in Central 1 region	0.0689	0.0179	3.8400	0.0000	0.0338	0.1041
Household lives in Central 2 region	-0.0360	0.0195	-1.8500	0.0640	-0.0742	0.0022
Household lives in Eastern region	0.0315	0.0195	1.6100	0.1080	-0.0069	0.0698
Household lives in San Salvador	0.1422	0.0208	6.8400	0.0000	0.1015	0.1829
Household lives in rural area	-0.0578	0.0162	-3.5600	0.0000	-0.0896	-0.0260
HH main cooking fuel is firewood	-0.3092	0.0172	-17.9700	0.0000	-0.3429	-0.2754
HH rents dwelling	0.1184	0.0230	5.1500	0.0000	0.0733	0.1635
Dwelling flooring material is ceramic brick	0.0907	0.0267	3.4000	0.0010	0.0384	0.1430
Household head worked in the last week	0.2130	0.0158	13.4500	0.0000	0.1820	0.2440
Number of TVs HH owns	0.1042	0.0146	7.1200	0.0000	0.0755	0.1329
Number of vehicles HH owns	0.3211	0.0193	16.6100	0.0000	0.2832	0.3590
Number of rooms used by HH	0.0710	0.0062	11.3600	0.0000	0.0588	0.0833
HH owns one or more blenders	0.1107	0.0167	6.6200	0.0000	0.0779	0.1434
HH owns one or more computers	0.2549	0.0294	8.6700	0.0000	0.1973	0.3125
HH owns one or more fans	0.0981	0.0163	6.0300	0.0000	0.0662	0.1300
HH owns one or more iron	0.1544	0.0182	8.4700	0.0000	0.1186	0.1901
HH owns one or more microwaves	0.0970	0.0248	3.9100	0.0000	0.0483	0.1456
HH owns one or more VCR or DVD players	0.1041	0.0162	6.4200	0.0000	0.0723	0.1359
HH owns one or more washing machines	0.1144	0.0276	4.1400	0.0000	0.0602	0.1685
HH disposes of garbage by public collection	0.0653	0.0176	3.7100	0.0000	0.0308	0.0997

## **Annex 1: Poverty Prediction at the National 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 El Salvador are defined as those whose predicted incomes fall below (above) the national line.

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 El Salvador, these five lines are the food poverty line, halfway between the median and national poverty line, the median 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**

<b>El Salvador</b> National Line Share of Poor: 39.9%	<b>Total Accuracy</b>	<b>Poverty Accuracy</b>	<b>Under-coverage</b>	<b>Leakage</b>	<b>PIE</b>	<b>BPAC</b>
<b>Single-step methods</b>						
Quantile regression (estimation point: 46)	75.82	69.60	30.40	29.90	-0.20	69.10

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**

	<b>Number of households identified as poor by the tool</b>	<b>Number of households identified as not poor by the tool</b>
<b>Number of “true” poor households (as determined by benchmark survey)</b>	3490 (27.9%)	1525 (12.2%)
<b>Number of “true” not poor households (as determined by benchmark survey)</b>	1499 (12.0%)	5992 (47.9%)

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

.46 Quantile regression  
 Min sum of deviations 5539.992

Number of obs = 12506  
 Pseudo R2 = 0.3473

<b>Variable</b>	<b>Coef.</b>	<b>Std. Err.</b>	<b>t</b>	<b>P&gt; t </b>	<b>[95% Conf. Interval]</b>	
Intercept	0.8377	0.0693	12.0900	0.0000	0.7019	0.9735
Household size	-0.3281	0.0111	-29.4800	0.0000	-0.3499	-0.3063
Household size squared	0.0188	0.0010	19.4300	0.0000	0.0169	0.0207
Household head age	0.0141	0.0027	5.3300	0.0000	0.0089	0.0193
Household head age squared	-0.0001	0.0000	-4.8100	0.0000	-0.0002	-0.0001
Household lives in Central 1 region	0.0824	0.0208	3.9600	0.0000	0.0416	0.1231
Household lives in Central 2 region	-0.0255	0.0225	-1.1300	0.2570	-0.0697	0.0186
Household lives in Eastern region	0.0321	0.0226	1.4200	0.1560	-0.0123	0.0765
Household lives in San Salvador	0.1307	0.0242	5.3900	0.0000	0.0832	0.1783
Household lives in rural area	-0.0449	0.0187	-2.4000	0.0160	-0.0815	-0.0083
HH main cooking fuel is firewood	-0.2831	0.0199	-14.2000	0.0000	-0.3222	-0.2440
HH rents dwelling	0.0903	0.0267	3.3800	0.0010	0.0379	0.1427
Dwelling flooring material is ceramic brick	0.0927	0.0310	2.9900	0.0030	0.0320	0.1534
Household head worked in the last week	0.2243	0.0183	12.2600	0.0000	0.1885	0.2602
Number of TVs HH owns	0.0945	0.0173	5.4500	0.0000	0.0605	0.1285
Number of vehicles HH owns	0.3032	0.0224	13.5400	0.0000	0.2593	0.3471
Number of rooms used by HH	0.0694	0.0073	9.5500	0.0000	0.0552	0.0837
HH owns one or more blenders	0.0916	0.0194	4.7300	0.0000	0.0537	0.1296
HH owns one or more computers	0.2583	0.0351	7.3700	0.0000	0.1896	0.3271
HH owns one or more fans	0.1130	0.0189	5.9900	0.0000	0.0760	0.1500
HH owns one or more iron	0.1541	0.0213	7.2400	0.0000	0.1124	0.1958
HH owns one or more microwaves	0.1122	0.0291	3.8500	0.0000	0.0551	0.1693
HH owns one or more VCR or DVD players	0.1067	0.0188	5.6700	0.0000	0.0698	0.1437
HH owns one or more washing machines	0.1468	0.0325	4.5200	0.0000	0.0832	0.2104
HH disposes of garbage by public collection	0.0721	0.0201	3.5800	0.0000	0.0327	0.1116

## 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 El Salvador tool was divided randomly into two data sets of 12,056 households (75 percent of the sample) and 4,168 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 shown in Table 7. The best model (1-step quantile) performs well in terms of BPAC and PIE, actually gaining in BPAC and losing 0.14 points for PIE, respectively.

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

	<b>Total Accuracy</b>	<b>Poverty Accuracy</b>	<b>Under-coverage</b>	<b>Leakage</b>	<b>PIE</b>	<b>BPAC</b>
<b>In-Sample Prediction</b>						
	81.10	52.84	47.16	46.83	-0.07	52.51
<b>Out-of-Sample Prediction</b>						
	81.78	56.04	43.96	45.01	0.21	54.99

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.<sup>5</sup> 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 +/- 7.9 percentage points. For PIE,

<sup>5</sup> This method of out-of-sample testing is used by Mark Schreiner for the PPI scorecards as detailed on [www.microfinance.com](http://www.microfinance.com)

which measures the difference between the predicted poverty rate and the actual poverty rate, the confidence interval is +/- 1.6 percentage points.

**Table 8: Bootstrapped Confidence Intervals on Assumption of Normality**

Variable	Mean	Std. Dev.	Confidence interval	
			LB	UB
<b>Total Accuracy</b>	82.11	0.78	80.59	83.64
<b>Poverty Accuracy</b>	54.61	2.11	50.47	58.74
<b>Undercoverage</b>	45.39	2.11	41.26	49.53
<b>Leakage</b>	43.04	3.21	36.74	49.33
<b>PIE</b>	-0.49	0.80	-2.06	1.07
<b>BPAC</b>	50.93	4.01	43.08	58.78

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<sup>th</sup> percentile of the sample distribution for each measure; the upper bound is defined by the 97.5<sup>th</sup> 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
<b>Total Accuracy</b>	80.56	83.61
<b>Poverty Accuracy</b>	50.22	58.54
<b>Undercoverage</b>	41.46	49.78
<b>Leakage</b>	37.13	49.58
<b>PIE</b>	-2.09	0.99
<b>BPAC</b>	41.51	56.86

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 El Salvador PAT appears to be slightly biased toward underestimating the actual extreme poverty rate, but nonetheless falls within a narrow band containing the true value in the population. By this measure, the predictive model behind the El Salvador PAT is accurate.