

Poverty Assessment Tool Accuracy Submission
USAID/IRIS Tool for Mexico
Submitted: July 19, 2010

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

1. Data source

For Mexico, existing data from the 2008 ENIGH were used to construct the poverty assessment tool. The full sample of 29,186 households is nationally representative. The sample used for tool construction comprises a randomly selected 21,650 households (75 percent of the full sample). The remainder, another randomly selected 7,536 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 Mexico, 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 Mexico 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 Mexico, the most accurate method, on the basis of BPAC, is the 1-step quantile regression.

Table 1: In-sample Accuracy Results for Prediction at the Legislative Poverty Line

Mexico Median line* Share of “very poor”: 20.1%	Total Accuracy	Poverty Accuracy	Under-coverage	Leakage	PIE	BPAC
Single-step methods						
OLS	83.28	40.46	59.54	23.27	-7.32	4.19
Quantile regression (estimation point: 37 percentile)	81.76	55.19	44.81	45.51	0.14	54.49
Linear Probability	83.15	27.85	72.15	11.32	-12.28	-32.98
Probit	83.72	37.97	62.03	18.61	-8.77	-5.45
Two-step methods						
OLS –99 percentile cutoff	83.28	40.48	59.52	23.31	-7.31	4.26
Quantile (estimation points: 37, 38) 99 percentile cutoff	81.90	54.78	45.22	44.43	-0.16	53.99
LP – 41 percentile cutoff	83.95	41.76	58.24	21.26	-7.47	4.78
Probit – 41 percentile cutoff	84.02	40.87	59.13	20.02	-7.90	1.76
*Median poverty line is 716 pesos per capita per month in 2008 prices in rural areas and 1,286 pesos in urban areas. This poverty line is based on the official national poverty line (patrimonio) of 1,282 in rural areas and 1,905 in urban areas.						

For Mexico, the functionality of predicting the poverty rate at other poverty lines—in this case, the food poverty line, the ‘capabilities’ line, the national line, and 50% above the national line —have 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.

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 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)³. The applicable poverty line for USAID tool development is the one that yields the higher household poverty rate for a given country.

There are three primary ‘national’ poverty lines in Mexico: food (alimentaria), capabilities (capacidades), and asset (patrimonio). The ‘capabilities’ line includes income to purchase a certain basket of food, along with a certain amount of income for health and education. The ‘asset’ line includes the elements of the capabilities line plus clothing, housing, and transport. The latter line is the closest to the typical concept of a food plus basic needs poverty line, and it was therefore selected as the national data line. The value of the line differs between urban and rural areas to account for price differences between the two. The national ‘asset’ line is 1,905 pesos per capita per month in urban areas and 1,282 in rural areas.

In Mexico, the median poverty line, or the household per capita income value of the 50th percentile below the national poverty line, is 1,286 pesos per capita per month in urban areas and 716 pesos 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 20.1% of households as “very poor.”

Alternatively, the international poverty line of \$1.25/day in 2005 PPP terms identifies 1.8% of households as “very poor.”⁴

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 Mexico is 7.648.

⁴ The World Bank’s PovcalNet provides a poverty headcount of 1.1% using population weights.

Hence the decision rule for Mexico’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 income 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 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

	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)	2,412 (11.1%)	1,959 (9.0%)
Number of “true” not very-poor households (as determined by benchmark survey)	1,989 (9.2%)	15,290 (70.7%)

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

.37 Quantile regression
 Min sum of deviations 9395.328

Number of obs = 21,650
 Pseudo R2 = 0.3672

Variable	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
Intercept	7.0644	0.0673	104.9300	0.0000	6.9324	7.1963
Household size	-0.3521	0.0094	-37.4500	0.0000	-0.3706	-0.3337
Household size squared	0.0190	0.0009	22.0300	0.0000	0.0173	0.0207
Household head age	0.0176	0.0024	7.2900	0.0000	0.0129	0.0223
Household head age squared	-0.0002	0.0000	-6.9400	0.0000	-0.0002	-0.0001
Household lives in rural area	-0.2144	0.0187	-11.4900	0.0000	-0.2509	-0.1778
HH lives in Coahuila de Zaragoza, Chihuahua, Sinaloa or Sonora	-0.0900	0.0213	-4.2300	0.0000	-0.1317	-0.0482
HH lives in Veracruz de Ignacio de la Llave	0.0308	0.0366	0.8400	0.4010	-0.0410	0.1026
HH lives in Campeche, Chiapas, Quintana Roo or Yucatán	-0.1036	0.0211	-4.9100	0.0000	-0.1450	-0.0622
HH lives in Guerrero or Oaxaca	-0.0061	0.0265	-0.2300	0.8170	-0.0580	0.0457
HH lives in Colima, Jalisco, Michoacán de Ocampo or Nayarit	-0.0699	0.0209	-3.3400	0.0010	-0.1108	-0.0289
HH lives in Baja California Sur	0.0303	0.0693	0.4400	0.6620	-0.1055	0.1661
HH lives in Baja California	-0.0226	0.0441	-0.5100	0.6080	-0.1089	0.0638
Dwelling floor is made of mosaic, marble, or tile	0.1438	0.0165	8.7200	0.0000	0.1115	0.1761
Drinking water is bought in jug or bottle	0.1329	0.0166	8.0100	0.0000	0.1004	0.1654
Garbage is disposed by burning	-0.1943	0.0250	-7.7800	0.0000	-0.2432	-0.1453
Number of rooms in dwelling	0.0635	0.0049	12.9300	0.0000	0.0539	0.0731
HH owns one or more TVs	0.2395	0.0295	8.1100	0.0000	0.1816	0.2974
HH owns one or more DVD or video disk players	0.1325	0.0156	8.4800	0.0000	0.1019	0.1631
HH owns one or more electric toasters	0.1612	0.0197	8.1900	0.0000	0.1226	0.1998
HH owns one or more microwaves	0.1302	0.0165	7.8700	0.0000	0.0978	0.1626
HH owns one or more refrigerators	0.1403	0.0211	6.6400	0.0000	0.0988	0.1817
HH owns one or more washing machines	0.1053	0.0157	6.7000	0.0000	0.0745	0.1361
HH owns one or more vacuum cleaners	0.1563	0.0262	5.9700	0.0000	0.1050	0.2077
HH owns one or more computers	0.2524	0.0189	13.3400	0.0000	0.2153	0.2894
HH owns one or more cars	0.2716	0.0173	15.6900	0.0000	0.2377	0.3055
HH owns one or more vans	0.2033	0.0217	9.3600	0.0000	0.1607	0.2459
HH owns one or more pick ups	0.1065	0.0221	4.8200	0.0000	0.0632	0.1499

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 Mexico 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 Mexico, these five lines are the food poverty line, capabilities line, median line, 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

Mexico National Line Share of Poor: 40.2%	Total Accuracy	Poverty Accuracy	Under-coverage	Leakage	PIE	BPAC
Single-step methods						
Quantile regression (estimation point: 45)	76.88	71.59	28.41	29.07	0.26	70.93

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)	6,235 (28.8%)	2,474 (11.4%)
Number of “true” not poor households (as determined by benchmark survey)	2,531 (11.7%)	10,410 (48.1%)

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

.45 Quantile regression
 Min sum of deviations 9801.167

Number of obs = 21,650
 Pseudo R2 = 0.3722

Variable	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
Intercept	7.1585	0.0637	112.3200	0.0000	7.0336	7.2834
Household size	-0.3528	0.0088	-40.1600	0.0000	-0.3701	-0.3356
Household size squared	0.0187	0.0008	23.6900	0.0000	0.0172	0.0203
Household head age	0.0185	0.0023	8.0800	0.0000	0.0140	0.0229
Household head age squared	-0.0002	0.0000	-7.7100	0.0000	-0.0002	-0.0001
Household lives in rural area	-0.1944	0.0177	-10.9900	0.0000	-0.2290	-0.1597
HH lives in Coahuila de Zaragoza, Chihuahua, Sinaloa or Sonora	-0.1084	0.0200	-5.4200	0.0000	-0.1476	-0.0692
HH lives in Veracruz de Ignacio de la Llave	0.0709	0.0342	2.0700	0.0380	0.0038	0.1380
HH lives in Campeche, Chiapas, Quintana Roo or Yucatán	-0.0721	0.0197	-3.6600	0.0000	-0.1107	-0.0335
HH lives in Guerrero or Oaxaca	0.0107	0.0253	0.4300	0.6710	-0.0388	0.0603
HH lives in Colima, Jalisco, Michoacán de Ocampo or Nayarit	-0.0640	0.0195	-3.2800	0.0010	-0.1022	-0.0258
HH lives in Baja California Sur	0.0839	0.0631	1.3300	0.1830	-0.0397	0.2076
HH lives in Baja California	-0.0294	0.0414	-0.7100	0.4770	-0.1105	0.0517
Dwelling floor is made of mosaic, marble, or tile	0.1482	0.0155	9.5500	0.0000	0.1178	0.1787
Drinking water is bought in jug or bottle	0.1250	0.0157	7.9800	0.0000	0.0943	0.1558
Garbage is disposed by burning	-0.1903	0.0233	-8.1700	0.0000	-0.2359	-0.1447
Number of rooms in dwelling	0.0642	0.0046	13.9700	0.0000	0.0552	0.0732
HH owns one or more TVs	0.2283	0.0275	8.3000	0.0000	0.1744	0.2822
HH owns one or more DVD or video disk players	0.1248	0.0148	8.4600	0.0000	0.0959	0.1537
HH owns one or more electric toasters	0.1666	0.0186	8.9300	0.0000	0.1300	0.2031
HH owns one or more microwaves	0.1209	0.0157	7.7200	0.0000	0.0902	0.1516
HH owns one or more refrigerators	0.1584	0.0199	7.9500	0.0000	0.1193	0.1974
HH owns one or more washing machines	0.1078	0.0147	7.3100	0.0000	0.0789	0.1367
HH owns one or more vacuum cleaners	0.1817	0.0244	7.4400	0.0000	0.1338	0.2295
HH owns one or more computers	0.2453	0.0179	13.6700	0.0000	0.2102	0.2805
HH owns one or more cars	0.2764	0.0163	16.9300	0.0000	0.2444	0.3084
HH owns one or more vans	0.1944	0.0207	9.3900	0.0000	0.1538	0.2350
HH owns one or more pick ups	0.1272	0.0208	6.1100	0.0000	0.0864	0.1681

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 Mexico tool was divided randomly into two data sets of 21,650 households (75 percent of the sample) and 7,536 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, losing 4.2 points for BPAC and losing 0.6 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						
	81.76	55.19	44.81	45.51	0.14	54.49
Out-of-Sample Prediction						
	82.02	53.73	46.27	42.80	-0.70	50.25

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 +/- 5.1 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 +/- 1.4 percentage points.

Table 8: Bootstrapped Confidence Intervals on Assumption of Normality

Variable	Mean	Std. Dev.	Confidence interval	
			LB	UB
Total Accuracy	81.28	0.65	80.01	82.55
Poverty Accuracy	54.00	1.86	50.35	57.65
Undercoverage	46.00	1.86	42.35	49.65
Leakage	48.04	3.02	42.11	53.97
PIE	0.39	0.72	-1.02	1.81
BPAC	50.66	2.59	45.59	55.73

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	80.01	82.65
Poverty Accuracy	50.37	57.62
Undercoverage	42.38	49.63
Leakage	42.42	54.36
PIE	-1.07	1.84
BPAC	44.88	55.03

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 Mexico PAT appears to be slightly biased toward underestimating the actual extreme poverty rate, but nonetheless will fall within 1.9 percentage points of the true value in the population (with 95 percent confidence). By this measure, the predictive model behind the Mexico PAT is accurate.