Poverty Assessment Tool Accuracy Submission
USAID/IRIS Tool for Paraguay
Submitted: March 30, 2009

The following report is divided into six sections. Section 1 provides a general overview of the tool development process. Section 2 describes the data set used to create the Poverty Assessment Tool for Paraguay. Section 3 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 4 reports on the in-sample accuracy of each prediction model considered. Sections 5 and 6 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 supplies a careful consideration of out-of-sample accuracy for the Paraguay Poverty Assessment Tool.

1. Overall approach to the tool development

The approach used to develop the Poverty Assessment Tool for Paraguay is built on the lessons learned and methods refined during the original USAID/IRIS project, Developing Poverty Assessment Tools (September 2003 to October 2006). In the initial phase of the project, the IRIS Center analyzed data from existing national household surveys and from surveys it conducted itself in countries where survey data were not available. The aim was to identify household indicators most closely associated with a household being "very poor" in terms of per-capita expenditures or income. IRIS used statistical methods to identify the 15 indicators that most closely track the per-capita expenditures or income of each household, as revealed by the household survey data. In addition, IRIS compared the performance of 8 different statistical methods in quantifying the statistical links between these 15 indicators and household expenditures/income; the accuracy of each method was assessed using criteria developed especially for this project. In this manner, IRIS identified the best-performing set of indicators (with associated weights) and statistical method for identifying the poverty status of households in each country. Statistical testing for accuracy was carried out for twelve countries in total.

In addition, the indicators that appeared among the "best 15" in at least one of the twelve countries were included in the next part of the project: testing for practicality. USAID selected seventeen microenterprise organizations to conduct field tests of practicality. Each question was rated as to whether the respondent found it to be sensitive, difficult, or that it was perceived that she falsified her answer. The lessons learned from the practicality testing were used to remove impractical indicators from consideration for the final poverty assessment tools.

The end result of this development process was a country-specific poverty assessment tool that estimates—rather than directly measures—household per capita consumption expenditure (or income) or the probability that a household is very poor based on a short set of practical

indicators. Each country tool is incorporated into a data entry template that allows a microenterprise practitioner to easily enter and store the responses of its sampled clients to indicator questions and also estimates the percentage of that practitioner's client households who are very poor.

In October 2006, USAID contracted the IRIS Center to build on the statistical methods and practicality information generated during the original project to develop poverty assessment tools for use by microenterprise practitioners in additional countries such as Paraguay. As part of this new phase, IRIS is exploring the use of existing household survey data sets beyond specifically the LSMS to develop new poverty assessment tools, refine the tool development and testing methodology, and strive to make the tools even simpler and easier to implement. IRIS is also providing a Help Desk to assist practitioners with the implementation of approved USAID poverty assessment tools.

2. Data source

For Paraguay, existing data from the 2000-2001 Enquesta Integrada de Hogares (EIH) were used to construct the poverty assessment tool. The full sample of 8,131 households is nationally representative. The sample used for tool construction comprises a randomly selected 4,066 households (50 percent of the full sample). The remainder, another randomly selected 4,065 households, is reserved for out-of-sample accuracy testing, which investigates the robustness of in-sample poverty estimation.

3. 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 Paraguay, more than 124 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]

4. Estimation methods used to identify final indicators and their weights/coefficients

As explained more fully in Section 6, the line used to construct the poverty tool for Paraguay is the "median poverty line" – the level of monthly expenditure that divides the poorest half of those living below the national poverty line from the less-poor half of the officially poor. Table 1 summarizes the accuracy results achieved by each of the eight estimation methods in predicting household poverty relative to this poverty line. For Paraguay, 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 (roughly two percentage points lower in BPAC and actually better in PIE) 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

PARAGUAY Median line***** Share of "very poor": 13.1%	**Total Accuracy**	**Poverty Accuracy**	**Under-coverage**	**Leakage**	**PIE**	**BPAC**
Single-step methods						
OLS	85.93	33.03	66.97	36.94	-4.07	3.01
Quantile regression (estimation point: 33)	**85.71**	**47.23**	**52.77**	**54.05**	**0.17**	**45.95**
Linear Probability	88.60	22.24	77.76	6.39	-9.67	-49.12
Probit	88.59	29.01	70.99	13.27	-7.82	-28.71
Two-step methods						
OLS – 48 percentile cutoff	87.87	42.77	57.23	32.35	-3.37	17.89
Quantile (estimation points: 33, 16) 48 percentile cutoff	87.36	51.33	48.67	45.77	-0.39	48.43
LP – 42 percentile cutoff	89.26	35.57	64.43	14.86	-6.71	-13.99
Probit – 42 percentile cutoff	88.74	33.21	66.79	16.33	-6.83	-17.24
* Median poverty line is 193,347 Paraguay Guarani per capita per month in 2001 prices in Asuncion, 199,368 in Central urban areas, 113,745 in other urban areas, and 55,948 in rural areas. This poverty line is based on the official national poverty line of 269,694 Paraguay Guarani per capita per month in Asuncion, 266,967 in Central urban areas, 166,762 in other urban areas, and 100,151 in rural areas.						

For Paraguay, the functionality of predicting the poverty rate at another poverty line—in this case, the national poverty line—has been added. When running the analysis routine with the Epi Info template, the user is presented the option to predict the extreme poverty rate (using the

[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

median line), the poverty rate (national line), or both. The methodology and the accuracy results for this prediction are discussed in Annex 1.

5. 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 EPI template used to calculate the incidence of extreme poverty among each implementing organization's clients. While a skilled EPI user would be able to locate the model's weights in the back-end, they would not be seen by the client or the interviewer during the normal course of interviewing, entering the data, or calculating the extreme poverty rate.

6. 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 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 Paraguay the applicable threshold is the median poverty line, the household per capita expenditure value of the 50th percentile below the national poverty line of 193,347 Paraguay Guarani per capita per month in Asuncion, 199,368 in Central urban areas, 113,745 in other urban areas, and 55,948 in rural areas, at the level of prices prevailing in 2000 and 2001 when the household survey data were collected.[4] At these values, the median poverty line identifies 13.1% of households as "very poor."

The alternative possibility for the poverty line is the "international poverty line" of $1.25/day in 2005 PPP terms. Expressed in prices prevailing at the time of data collection, the international poverty line is 49,895 per capita per month for Paraguay. This line identifies 6.7% of households in the sample as very poor.

Hence the decision rule for Paraguay's USAID poverty assessment tool in classifying the "very poor" (and the "not very-poor") is whether that predicted per capita monthly expenditures of a household fall below (or above) the median 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.

[4] These median poverty lines are based on the following national poverty lines (in 2000/01 prices): 269,694 Paraguay Guarani per capita per month in Asuncion, 266,967 in Central urban areas, 166,762 in other urban areas, and 100,151 in rural area. The national poverty lines yield a household poverty rate of 26.3 percent.

Because the selected tool is based on a Quantile model, each household whose estimated per capita consumption expenditures according to the tool fall below 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."

An additional requirement for using the median poverty line is that the national poverty line on which it depends is actively used by the local government. This appears to be the case in Paraguay, where the government uses the national poverty line for poverty monitoring.[5]

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)	243 (6.0%)	293 (7.2%)
Number of "true" not very-poor households (as determined by benchmark survey)	256 (6.3%)	3,274 (80.5%)

[5] See http://www.undp.org.py/enrepd/download/enrepdinforme.pdf

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

PARAGUAY 1 STEP MAXR/QUANT: variables from MAXR/OLS 100 percentile model
Regression results, estimation point of 33 percentile
NOTE: rural dropped due to collinearity

.33 Quantile regression Number of obs = 4066
Raw sum of deviations 3221.662 (about 12.059275)
Min sum of deviations 1950.883 Pseudo R2 = 0.3944

Variable	Coef.	Std. Err.	t	P>|t|	[95% Conf. Interval]	
Intercept	13.4134	0.1637	81.9500	0.0000	13.0925	13.7343
Household lives in Asunción	0.2138	0.0486	4.4000	0.0000	0.1186	0.3090
Household lives in Central urban area	0.1322	0.0551	2.4000	0.0160	0.0243	0.2401
Household lives in rural area	-0.1535	0.0512	-3.0000	0.0030	-0.2539	-0.0531
Household size	-0.2783	0.0187	-14.9200	0.0000	-0.3148	-0.2417
Household size squared	0.0110	0.0014	7.7700	0.0000	0.0082	0.0138
Household head age	-0.0141	0.0064	-2.2100	0.0270	-0.0267	-0.0016
Household head age squared	0.0001	0.0001	2.2300	0.0260	0.0000	0.0003
Household head's marital status is widowed	0.0366	0.0651	0.5600	0.5730	-0.0909	0.1642
Share of household members that are literate	0.4463	0.0931	4.7900	0.0000	0.2638	0.6288
Number of rooms in dwelling	0.0862	0.0130	6.6300	0.0000	0.0607	0.1117
Floor of dwelling is made of earth	-0.1846	0.0525	-3.5200	0.0000	-0.2876	-0.0817
Roof of dwelling is made of eternit	-0.1841	0.0486	-3.7900	0.0000	-0.2793	-0.0888
Roof of dwelling is made of straw or wood	-0.4571	0.0543	-8.4200	0.0000	-0.5634	-0.3507
Dwelling's water supply is located inside the dwelling	-0.1675	0.0568	-2.9500	0.0030	-0.2788	-0.0561
Dwelling's water supply is located on the property	-0.2400	0.0469	-5.1200	0.0000	-0.3319	-0.1481
Dwelling's water supply is located outside the property	-0.3918	0.0682	-5.7500	0.0000	-0.5254	-0.2581
Main cooking fuel is firewood	-0.3077	0.0546	-5.6400	0.0000	-0.4146	-0.2007
Household owns one or more VCRs	0.2676	0.0546	4.9000	0.0000	0.1605	0.3746
Household owns one or more air conditioners	0.3376	0.0665	5.0800	0.0000	0.2072	0.4679
Household owns one or more cars	0.2788	0.0507	5.5000	0.0000	0.1794	0.3783
Household owns one or more food processors	0.2211	0.0717	3.0800	0.0020	0.0806	0.3616
Number of chickens owned by household	0.0027	0.0008	3.4500	0.0010	0.0012	0.0042

Annex 1: Poverty Prediction at the National Poverty Line

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/day PPP models, this will be the $2/day PPP; for median poverty models, the "poor" threshold will be the national poverty line. Following this logic, then, the "poor" ("not poor") in Paraguay are defined as those whose predicted incomes fall below (above) the national poverty 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.

Table 4: Accuracy Results Obtained for Prediction at the National Poverty Line

Paraguay National Line Share of Poor: 26.3%	**Total Accuracy**	**Poverty Accuracy**	**Under-coverage**	**Leakage**	**PIE**	**BPAC**
Single-step methods						
Quantile regression (estimation point: 43)	79.83	63.34	36.66	39.62	0.78	60.38

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 Poverty 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)	617 (15.2%)	420 (10.3%)
Number of "true" not poor households (as determined by benchmark survey)	350 (8.6%)	2,679 (65.9%)

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

PARAGUAY 1 STEP MAXR/QUANT: variables from MAXR/OLS 100 percentile model
Regression results, estimation point of 43 percentile
NOTE: rural dropped due to collinearity

.43 Quantile regression Number of obs = 4066
Raw sum of deviations 3431.652 (about 12.351055)
Min sum of deviations 2056.744 Pseudo R2 = 0.4007

Variable	Coef.	Std. Err.	t	P>|t|	[95% Conf. Interval]	
Intercept	13.6369	0.1075	126.9100	0.0000	13.4262	13.8475
Household lives in Asunción	0.1965	0.0332	5.9200	0.0000	0.1314	0.2615
Household lives in Central urban area	0.0916	0.0372	2.4600	0.0140	0.0187	0.1646
Household lives in rural area	-0.1510	0.0336	-4.5000	0.0000	-0.2168	-0.0851
Household size	-0.2825	0.0143	-19.7800	0.0000	-0.3105	-0.2545
Household size squared	0.0115	0.0011	10.2100	0.0000	0.0093	0.0137
Household head age	-0.0155	0.0043	-3.6400	0.0000	-0.0239	-0.0072
Household head age squared	0.0001	0.0000	3.3000	0.0010	0.0001	0.0002
Household head's marital status is widowed	0.0500	0.0415	1.2000	0.2290	-0.0314	0.1313
Share of household members that are literate	0.3981	0.0602	6.6100	0.0000	0.2801	0.5161
Number of rooms in dwelling	0.0898	0.0089	10.0700	0.0000	0.0723	0.1073
Floor of dwelling is made of earth	-0.1671	0.0388	-4.3100	0.0000	-0.2431	-0.0911
Roof of dwelling is made of eternit	-0.2401	0.0310	-7.7600	0.0000	-0.3008	-0.1794
Roof of dwelling is made of straw or wood	-0.4660	0.0381	-12.2200	0.0000	-0.5407	-0.3912
Dwelling's water supply is located inside the dwelling	-0.1792	0.0399	-4.4900	0.0000	-0.2574	-0.1010
Dwelling's water supply is located on the property	-0.2030	0.0314	-6.4600	0.0000	-0.2647	-0.1414
Dwelling's water supply is located outside the property	-0.4119	0.0454	-9.0700	0.0000	-0.5009	-0.3229
Main cooking fuel is firewood	-0.3141	0.0371	-8.4600	0.0000	-0.3868	-0.2413
Household owns one or more VCRs	0.2764	0.0377	7.3400	0.0000	0.2026	0.3502
Household owns one or more air conditioners	0.3266	0.0362	9.0100	0.0000	0.2555	0.3976
Household owns one or more cars	0.2471	0.0341	7.2500	0.0000	0.1803	0.3139
Household owns one or more food processors	0.2526	0.0464	5.4500	0.0000	0.1617	0.3435
Number of chickens owned by household	0.0028	0.0006	4.9900	0.0000	0.0017	0.0039

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 2, the data set used to construct the Paraguay tool was divided randomly into two data sets of equal size (4,066 and 4,065 households). 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 only about 8 points and 1 point, 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						
	85.71	47.23	52.77	54.05	0.17	45.95
Out-of-Sample Prediction						
	85.51	48.54	51.46	61.76	1.32	38.24

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.

Table 8 presents the out-of-sample, bootstrapped confidence intervals for the 1-step Quantile model. The performance of this model is modest. The confidence interval around the sample mean BPAC is relatively wide at +/- 16.4 percentage points. For PIE, which measures the difference between the actual poverty rate in the sample and the predicted poverty rate, the confidence interval is +/- 2.5 percentage points.

[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: Bootstrapped Confidence Intervals Computed on Assumption of Normality

Accuracy Measure	Mean	Std. Err.	95% Confidence Interval	
			LB	UB
Total Accuracy	85.45	1.07	83.35	87.56
Poverty Accuracy	48.58	3.28	42.15	55.01
Undercoverage	51.42	3.28	44.99	57.85
Leakage	62.45	9.41	44.01	80.90
PIE	1.37	1.25	-1.08	3.81
BPAC	36.40	8.35	20.04	52.76

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^{th} percentile of the sample distribution for each measure; the upper bound is defined by the 97.5^{th} 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	83.24	87.36
Poverty Accuracy	42.35	55.02
Undercoverage	44.98	57.65
Leakage	46.95	83.15
PIE	-0.97	3.93
BPAC	16.85	48.49

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 Paraguay PAT appears to be somewhat biased toward overestimating the actual extreme poverty rate, but nonetheless will fall within at least four percentage points of the true value in the population (with greater than 95-percent confidence). By this measure, the predictive model behind the Paraguay PAT is moderately accurate.