

Poverty Assessment Tool Accuracy Submission
USAID/IRIS Tool for Bolivia
Revised to address median household issue: October 1, 2009

The following report is divided into five sections. Section 1 describes the data set used to create the Poverty Assessment Tool for Bolivia. 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 supplies a careful consideration of out-of-sample accuracy for the Bolivia Poverty Assessment Tool.

1. Data source

For Bolivia, existing data from the 2005 Encuesta de Hogares (EH) integrated survey were used to construct the poverty assessment tool. The full sample of 4,086 households is nationally representative. The sample used for tool construction comprises a randomly selected 2,043 households (50 percent of the full sample). The remainder, another randomly selected 2,043 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 Bolivia, more than 125 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 Bolivia is the “median poverty line” – the level of 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 Bolivia, 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

BOLIVIA Median line* Share of “very poor”: 24.2%	Total Accuracy	Poverty Accuracy	Under-coverage	Leakage	PIE	BPAC
Single-step methods						
OLS	84.36	58.51	41.49	21.66	-4.91	38.68
Quantile regression (estimation point: 42)	83.69	67.23	32.77	33.09	0.08	66.91
Linear Probability	83.62	49.77	50.23	15.92	-8.50	15.46
Probit	84.68	57.75	42.25	19.61	-5.61	35.12
Two-step methods						
OLS –46 percentile cutoff	84.61	60.26	39.74	22.41	-4.29	42.93
Quantile (estimation points: 42, 20) 46 percentile cutoff	84.04	68.00	32.00	32.51	0.13	67.49
LP – 50 percentile cutoff	85.13	63.02	36.98	23.07	-3.44	49.12
Probit –50 percentile cutoff	84.53	57.19	42.81	19.64	-5.74	34.03
*Poverty lines vary by department and, for El Alto only, by city. See Section 5 for details.						

For Bolivia, 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

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.

poverty rate (using the median line), the poverty rate (national line), or both. 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 Epi Info 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.

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 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 Bolivia the applicable threshold is the median poverty line, the household per capita expenditure value of the 50th percentile below the national poverty line, at the level of prices prevailing in 2005 when the household survey data were collected. In Bolivia, the official poverty line varies by department (and also for El Alto, a suburb of La Paz) to account for spatial price differences (see table below). At these values, the median poverty line identifies 24.2 % of households as “very poor.” The 2005 EH data set includes a household income aggregate in addition to the household expenditure aggregate. The Bolivian government and certain external sources (including the World Bank on PovcalNet) use the household income aggregate, despite its well- known limitations.

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 Bolivia is 2.57.

Official National Poverty Lines and Median Extreme Poverty Lines for Bolivia
Monthly expenditures per capita in Bolivianos

Location Variables	National Poverty Line	Median Poverty Line
Chuquisaca	367.32	220.17
El Alto	299.56	228.98
La Paz	368.77	257.29
Cochabamba	384.50	265.97
Oruro	335.34	249.13
Potosí	308.42	190.57
Tarija	384.50	275.90
Santacruz	388.83	269.02
Beni	388.83	248.35
Pando	388.83	373.82
Rural	281.52	149.00

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 97.81 bolivianos per capita per month for Bolivia.⁴ This line identifies 6.3% of households in the sample as very poor.⁵

Hence the decision rule for Bolivia’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 is less than or equal to (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.”

⁴ Despite repeated attempts to contact the national statistical office in Bolivia, it is not clear whether expenditure values collected in the data set are adjusted for inflation to a certain time period (e.g., a specific month, or average 2005 prices). We assume average 2005 prices. Therefore, the \$1.25 PPP extreme poverty line equals $1.25 \times 2.57 \times (365.25/12)$ or 97.81 bolivianos per capita per month. Even if the data values required large adjustments for inflation in 2005, the PPP line would not eclipse the median lines.

⁵ As mentioned, World Bank’s PovcalNet estimates the distribution of household income, rather than expenditures, to yield a poverty rate of 19.6% at the international \$1.25 PPP line. We obtain a result of 22.0% in our own tests using the 2005 EH income data.

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 Bolivia, where the government uses the national poverty line for poverty monitoring.⁶

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)	340 (16.6%)	166 (8.1%)
Number of “true” not very-poor households (as determined by benchmark survey)	167 (8.2%)	1370 (67.1%)

⁶ This is true as of the 2005 World Bank Poverty Assessment for Bolivia: <http://go.worldbank.org/PVQ19004U0>

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

.42 Quantile regression

Number of obs = 2043

Min sum of deviations 691.9926

Pseudo R2 = 0.4857

Variable	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
Intercept	6.0427	0.1040	58.1000	0.0000	5.8388	6.2467
Household size	-0.3713	0.0164	-22.6200	0.0000	-0.4034	-0.3391
Household size squared	0.0177	0.0015	12.0000	0.0000	0.0148	0.0206
Household head age	0.0178	0.0040	4.4300	0.0000	0.0099	0.0257
Household head age squared	-0.0002	0.0000	-4.9400	0.0000	-0.0003	-0.0001
Household lives in Cochabamba	0.0020	0.0349	0.0600	0.9550	-0.0664	0.0704
Household lives in Oruro	0.0294	0.0449	0.6600	0.5120	-0.0586	0.1173
Household lives in Potosí	-0.2436	0.0381	-6.3900	0.0000	-0.3183	-0.1689
Household lives in Tarija	0.0839	0.0413	2.0300	0.0430	0.0028	0.1649
Household lives in Santa Cruz	0.0631	0.0378	1.6700	0.0960	-0.0111	0.1373
Household lives in Pando	0.1792	0.0809	2.2100	0.0270	0.0205	0.3379
Household lives in Chuquisaca	-0.1948	0.0427	-4.5600	0.0000	-0.2786	-0.1110
Household lives in Beni	0.2518	0.0522	4.8300	0.0000	0.1495	0.3541
Household lives in rural area	-0.1560	0.0322	-4.8400	0.0000	-0.2191	-0.0928
Dwelling is subleased	0.1807	0.0327	5.5200	0.0000	0.1165	0.2448
Wall of dwelling is made of brick	0.1602	0.0294	5.4400	0.0000	0.1024	0.2179
Wall of dwelling is made of wood	0.3193	0.0616	5.1800	0.0000	0.1984	0.4402
Floor of dwelling is made of dirt	-0.2117	0.0393	-5.3900	0.0000	-0.2888	-0.1347
Floor of dwelling is made of cement	-0.1386	0.0284	-4.8800	0.0000	-0.1942	-0.0829
Household owns one or more refrigerators	0.1925	0.0318	6.0600	0.0000	0.1301	0.2548
Household owns one or more radio-cassette players	0.0767	0.0225	3.4100	0.0010	0.0325	0.1209
Household owns one or more televisions	0.1281	0.0348	3.6800	0.0000	0.0599	0.1963
Household owns one or more VCRs or DVD players	0.2309	0.0305	7.5700	0.0000	0.1711	0.2907
Household owns one or more fans	0.1077	0.0423	2.5500	0.0110	0.0248	0.1906
Household owns one or more cars	0.1625	0.0499	3.2600	0.0010	0.0646	0.2605
Number of beds owned	0.0725	0.0091	7.9500	0.0000	0.0546	0.0904
Number of kitchens in dwelling	0.1153	0.0245	4.7000	0.0000	0.0673	0.1634
Number of computers owned	0.3994	0.0390	10.2400	0.0000	0.3229	0.4758
Household owns one or more sheep	-0.1254	0.0366	-3.4200	0.0010	-0.1972	-0.0535

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 \$1.25/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 Bolivia 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

Bolivia National Line Share of Poor: 48.3%	Total Accuracy	Poverty Accuracy	Under-coverage	Leakage	PIE	BPAC
Single-step methods						
Quantile regression (estimation point: 47)	80.40	80.16	19.84	21.18	0.64	78.82

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)	782 (38.3%)	194 (9.5%)
Number of “true” not poor households (as determined by benchmark survey)	207 (10.1%)	860 (42.1%)

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

.47 Quantile regression

Number of obs = 2043

Min sum of deviations 705.1879

Pseudo R2 = 0.4838

Variable	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
Intercept	6.1393	0.1278	48.0400	0.0000	5.8887	6.3899
Household size	-0.3697	0.0201	-18.3800	0.0000	-0.4092	-0.3303
Household size squared	0.0174	0.0018	9.6100	0.0000	0.0138	0.0209
Household head age	0.0160	0.0049	3.2400	0.0010	0.0063	0.0257
Household head age squared	-0.0002	0.0001	-3.7400	0.0000	-0.0003	-0.0001
Household lives in Cochabamba	0.0215	0.0429	0.5000	0.6160	-0.0627	0.1058
Household lives in Oruro	0.0285	0.0541	0.5300	0.5990	-0.0776	0.1345
Household lives in Potosí	-0.2240	0.0464	-4.8300	0.0000	-0.3150	-0.1330
Household lives in Tarija	0.0520	0.0511	1.0200	0.3090	-0.0483	0.1523
Household lives in Santa Cruz	0.0500	0.0464	1.0800	0.2820	-0.0410	0.1409
Household lives in Pando	0.1792	0.0991	1.8100	0.0710	-0.0152	0.3735
Household lives in Chuquisaca	-0.1753	0.0531	-3.3000	0.0010	-0.2793	-0.0712
Household lives in Beni	0.2729	0.0638	4.2800	0.0000	0.1478	0.3981
Household lives in rural area	-0.1438	0.0387	-3.7100	0.0000	-0.2198	-0.0678
Dwelling is subleased	0.1702	0.0401	4.2500	0.0000	0.0916	0.2487
Wall of dwelling is made of brick	0.1772	0.0360	4.9200	0.0000	0.1065	0.2478
Wall of dwelling is made of wood	0.3132	0.0741	4.2300	0.0000	0.1679	0.4585
Floor of dwelling is made of dirt	-0.2154	0.0479	-4.4900	0.0000	-0.3094	-0.1214
Floor of dwelling is made of cement	-0.1572	0.0349	-4.5100	0.0000	-0.2257	-0.0888
Household owns one or more refrigerators	0.1866	0.0393	4.7500	0.0000	0.1096	0.2636
Household owns one or more radio-cassette players	0.0874	0.0278	3.1500	0.0020	0.0330	0.1419
Household owns one or more televisions	0.1272	0.0430	2.9600	0.0030	0.0429	0.2115
Household owns one or more VCRs or DVD players	0.2145	0.0374	5.7300	0.0000	0.1410	0.2879
Household owns one or more fans	0.1291	0.0524	2.4600	0.0140	0.0264	0.2318
Household owns one or more cars	0.1958	0.0606	3.2300	0.0010	0.0770	0.3146
Number of beds owned	0.0706	0.0112	6.3200	0.0000	0.0487	0.0924
Number of kitchens in dwelling	0.1185	0.0295	4.0200	0.0000	0.0607	0.1763
Number of computers owned	0.4056	0.0479	8.4700	0.0000	0.3116	0.4996
Household owns one or more sheep	-0.1327	0.0446	-2.9700	0.0030	-0.2202	-0.0451

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 Bolivia tool was divided randomly into two data sets of equal size (2,043 and 2,043 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 moderately well in terms of BPAC and PIE, losing about 17 points and 1.9 points, 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						
	83.65	67.18	32.82	33.29	0.11	66.71
Out-of-Sample Prediction						
	81.70	57.22	42.78	35.33	-1.74	49.78

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.⁷ 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 good. The confidence interval around the sample mean BPAC is relatively narrow at +/- 12.0 percentage points. For PIE, which

⁷ This method of out-of-sample testing is used by Mark Schreiner for the PPI scorecards as detailed on www.microfinance.com

measures the difference between the predicted poverty rate and the actual poverty rate, the confidence interval is +/- 2.2 percentage points.

Table 8: Bootstrapped Confidence Intervals Computed on Assumption of Normality

Accuracy Measure	Mean	Std. Err.	95% Confidence Interval	
			LB	UB
Total Accuracy	81.88	1.02	79.89	83.87
Poverty Accuracy	57.58	2.61	52.47	62.69
Undercoverage	42.42	2.61	37.30	47.53
Leakage	34.30	3.60	27.24	41.37
PIE	-1.94	1.10	-4.10	0.21
BPAC	49.33	6.11	37.36	61.31

The results presented in Table 8 assume a normal distribution for the accuracy measures from the bootstrapped samples. This ignores the possibility that these estimates may have a skewed distribution. Table 9 presents alternative 95% confidence intervals. The lower bound is defined by the 2.5th percentile of the sample distribution for each measure; the upper bound is defined by the 97.5th percentile. On the whole, the results are quite similar between Tables 8 and 9.

Table 9: Bootstrapped Confidence Intervals Computed Empirically from Sampling Distribution without Normality Assumption

Accuracy Measure	95% Confidence Interval	
	LB	UB
Total Accuracy	79.78	83.68
Poverty Accuracy	52.51	62.65
Undercoverage	37.35	47.49
Leakage	27.60	41.66
PIE	-4.06	0.13
BPAC	36.73	60.48

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 Bolivia PAT appears to be somewhat biased toward underestimating the actual extreme poverty rate, but nonetheless will fall within 4.3 percentage points of the true value in the population (with greater than 95-percent confidence). By this measure, the predictive model behind the Bolivia PAT is accurate.