Poverty Assessment Tool Accuracy Submission
USAID/IRIS Tool for Cambodia
Submitted: June 27, 2008
Revised with new 2005 PPP: September 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 Poverty Assessment Tool for Cambodia. 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 additional poverty lines beyond that required by the Congressional legislation. Annex 2 supplies a careful consideration of out-of-sample accuracy for the Cambodia Poverty Assessment Tool.

1. Overall approach to the tool development

The approach used to develop the poverty assessment tool for Cambodia 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 will also estimate 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 Cambodia. As part of this new phase, IRIS will explore 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 will also provide a Help Desk to assist practitioners with the implementation of approved USAID poverty assessment tools.

2. Data source

For Cambodia, existing data from the 2004 Cambodia Socioeconomic Survey (CSES) were used to construct the poverty assessment tool. The full sample of 14,984 households is nationally representative. The sample used for tool construction is comprised of a randomly selected 11,238 households (75 percent of the full sample). The remainder, another 3,746 randomly selected 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 Cambodia, more than 100 indicators from all categories, except for illness and disability and employment, 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 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 Cambodia is the $1.25/day PPP line. Table 1 summarizes the accuracy results achieved by each of the eight estimation methods in predicting household poverty relative to this line. For Cambodia, the most accurate method, on the basis of BPAC, is the 2-step Linear Probability regression. However, the 1-step Quantile regression is slightly less accurate and requires only 15 indicators. Following precedent, 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 Lines

Cambodia $1.25/day PPP line***** Share of "very poor": 35.2%	**Total Accuracy**	**Poverty Accuracy**	**Under-coverage**	**Leakage**	**PIE**	**BPAC**
Single-step methods						
OLS	76.44	58.23	41.77	24.57	-6.11	41.03
Quantile regression (estimation point: 46)	**76.54**	**66.75**	**33.25**	**32.81**	**-0.16**	**66.31**
Linear Probability	76.65	60.85	39.15	26.59	-4.46	48.28
Probit	76.45	60.60	39.40	26.91	-4.43	48.11
Two-step methods						
OLS – 48 percentile cutoff	77.36	62.19	37.81	25.95	-4.21	50.33
Quantile (estimation points: 46, 23) 48 percentile cutoff	77.13	68.27	31.73	32.65	0.33	67.34
LP – 45 percentile cutoff	77.12	67.73	32.27	32.14	-0.05	67.59
Probit – 45 percentile cutoff	77.01	65.82	34.18	30.56	-1.29	62.21
* The $1.25/day international poverty line is 1,899 riel per capita per day in average 2004 prices.						

For Cambodia, the functionality of predicting the poverty rate at another poverty line—in this case, the $2.50/day PPP—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 $1.25/line), the poverty rate (using the $2.50 line), or both. 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

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 backend, 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 lines established by the national government or those living on the local equivalent of less than $1.25/day per capita (in Purchasing Power Parity, or PPP, terms). The applicable poverty line for USAID tool development is the one that yields the higher household poverty rate for a given country.

In Cambodia, the median poverty lines are defined as the household per capita expenditure value of the 50th percentile below the national poverty lines, at the level of average prices prevailing in 2004 when the household survey data were collected. In Cambodia there are three official national poverty lines: Phnom Penh (2,351 riels per capita per day), other urban areas (1,952), and rural areas (1,753). These national lines yielded median poverty lines of 1819, 1531, and 1349 riels respectively. Applying these median poverty lines identifies 15.3% of households as "very poor".

Alternatively, the $1.25/day international poverty line is 1,899 riel per capita per day. [3] It yields a household poverty rate of 35.2%.[4] Hence the decision rule for Cambodia's USAID poverty assessment tool in classifying the "very poor" (and the "not very-poor") is whether the predicted per capita monthly expenditures of a household fall below (or above) the international $1.25/day poverty line.

Because the selected tool is based on a Quantile model, each household whose estimated per capita consumption expenditures is equal to or below $1.25/day poverty line, depending on location, is identified as "very poor," and each household whose estimated per capita consumption expenditures exceeds this line is identified as "not very-poor."

---

[3] The relevant 2005 PPP value for Cambodia is 1615.2985 in average 2005 prices. The PPP poverty line is therefore equal to 1.25*1615.2985*(94.03/100), where the last two terms are the average monthly CPI in 2004 and 2005 respectively.

[4] PovcalNet reports a *population* headcount of 40.2%. We find a population headcount of 40.1% using our data.

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)	2664 (23.7%)	1327 (11.8%)
Number of "true" not very-poor households (as determined by benchmark survey)	1309 (11.6%)	5938 (52.9%)

Table 3: Regression Estimates using 1-step Quantile Method for Prediction at the $1.25/day Poverty Line

.46 Quantile regression — Number of obs = 11238
Min sum of deviations 3792.145 — Pseudo R2 = 0.3415

Variable	Coef.	Std. Err.	t	P>|t|	[95% Conf. Interval]	
Intercept	8.0189	0.0445	180.0000	0.0000	7.9316	8.1063
Household head age	0.0152	0.0018	8.2800	0.0000	0.0116	0.0188
Household head age squared	-0.0001	0.0000	-7.7700	0.0000	-0.0002	-0.0001
Household size	-0.2810	0.0081	-34.6700	0.0000	-0.2969	-0.2652
Household size squared	0.0143	0.0007	20.6500	0.0000	0.0129	0.0157
Household lives in Phnom Penh	0.2306	0.0199	11.6000	0.0000	0.1916	0.2695
Household lives on the plains region	0.0352	0.0095	3.6900	0.0000	0.0165	0.0539
Household lives on the coastal region	0.0790	0.0153	5.1500	0.0000	0.0489	0.1091
Household lives on the plateau/mountain region	-0.0666	0.0137	-4.8700	0.0000	-0.0934	-0.0398
Household lives in rural area	-0.0807	0.0121	-6.7000	0.0000	-0.1044	-0.0571
Household head is female	-0.1252	0.0100	-12.4600	0.0000	-0.1449	-0.1055
Share of literate household members over 5 years old	0.2653	0.0217	12.2400	0.0000	0.2228	0.3078
Roof is made of thatch	-0.0732	0.0101	-7.2300	0.0000	-0.0931	-0.0534
Main source of lighting is publicly-provided electricity	0.3503	0.0181	19.3600	0.0000	0.3148	0.3857
Main source of lighting is privately-generated electricity	0.2953	0.0177	16.7100	0.0000	0.2607	0.3299
Main source of lighting is battery	0.1011	0.0104	9.7100	0.0000	0.0807	0.1215
Household treats its drinking water	0.1026	0.0089	11.5700	0.0000	0.0852	0.1200
Main cooking fuel is liquefied petroleum gas	0.3033	0.0202	15.0200	0.0000	0.2637	0.3429
Household owns one or more televisions	0.1095	0.0095	11.4900	0.0000	0.0908	0.1282
Household owns one or more video tape players or recorders	0.1518	0.0173	8.7900	0.0000	0.1180	0.1857
Household owns one or more motorcycles	0.1402	0.0099	14.1600	0.0000	0.1208	0.1596
Household owns one or more suitcases	0.1087	0.0095	11.4100	0.0000	0.0900	0.1274
Household owns one or more dining sets	0.1986	0.0165	12.0400	0.0000	0.1662	0.2309
Household owns one or more wardrobes or cabinets	0.1589	0.0121	13.1200	0.0000	0.1352	0.1826
Household has consumed meat in the past 7 days	0.0242	0.0015	15.7200	0.0000	0.0211	0.0272

Annex 1: Poverty Prediction at the $2.50/day PPP 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 lines. Following this logic, then, the "poor" in Cambodia are defined as those whose predicted incomes is less than or equal to the $2.50/day PPP line.
.

Table 4 summarizes the predictive accuracy results for the $2.50/day line using the Quantile model specification from the $1.25/day line. The indicators are the same as those in the model for the $1.25/day 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 $2.50/day line.

Table 4: Accuracy Results Obtained for Prediction at the $2.50/day PPP Line

Cambodia $2.50/day PPP Line Share of Poor: 75.7%	**Total Accuracy**	**Poverty Accuracy**	**Under-coverage**	**Leakage**	**PIE**	**BPAC**
Single-step methods						
Quantile regression (estimation point: 65)	84.50	89.90	10.10	10.37	0.21	89.63

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 $2.50/day 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)	7647 (68.0%)	859 (7.6%)
Number of "true" not poor households (as determined by benchmark survey)	882 (7.8%)	1850 (16.6%)

Table 6: Regression Estimates using 1-step Quantile Method for Prediction at the $2.50/day Line

.65 Quantile regression Number of obs = 11238
Min sum of deviations 3680.977 Pseudo R2 = 0.3626

Variable	Coef.	Std. Err.	t	P>|t|	[95% Conf. Interval]	
Intercept	8.2403	0.0603	136.7000	0.0000	8.1222	8.3585
Household head age	0.0134	0.0025	5.3500	0.0000	0.0085	0.0183
Household head age squared	-0.2892	0.0109	-26.6300	0.0000	-0.3105	-0.2679
Household size	-0.0001	0.0000	-5.0000	0.0000	-0.0002	-0.0001
Household size squared	0.0150	0.0009	16.3600	0.0000	0.0132	0.0168
Household lives in Phnom Penh	0.2174	0.0273	7.9700	0.0000	0.1639	0.2708
Household lives on the plains region	0.0232	0.0128	1.8200	0.0690	-0.0018	0.0482
Household lives on the coastal region	0.0795	0.0207	3.8400	0.0000	0.0389	0.1201
Household lives on the plateau/mountain region	-0.0574	0.0184	-3.1200	0.0020	-0.0934	-0.0214
Household lives in rural area	-0.0756	0.0160	-4.7200	0.0000	-0.1069	-0.0442
Household head is female	-0.1171	0.0135	-8.6900	0.0000	-0.1435	-0.0907
Share of literate household members over 5 years old	0.3221	0.0291	11.0700	0.0000	0.2650	0.3792
Roof is made of thatch	-0.0679	0.0136	-5.0000	0.0000	-0.0945	-0.0413
Main source of lighting is publicly-provided electricity	0.3863	0.0240	16.1000	0.0000	0.3393	0.4334
Main source of lighting is privately-generated electricity	0.2889	0.0238	12.1600	0.0000	0.2424	0.3355
Main source of lighting is battery	0.1133	0.0140	8.1200	0.0000	0.0860	0.1407
Household treats its drinking water	0.1193	0.0119	10.0600	0.0000	0.0961	0.1426
Main cooking fuel is liquefied petroleum gas	0.2893	0.0283	10.2400	0.0000	0.2339	0.3447
Household owns one or more televisions	0.1082	0.0127	8.5000	0.0000	0.0833	0.1332
Household owns one or more video tape players or recorders	0.1776	0.0233	7.6300	0.0000	0.1320	0.2233
Household owns one or more motorcycles	0.1434	0.0134	10.7000	0.0000	0.1171	0.1697
Household owns one or more suitcases	0.1223	0.0128	9.5600	0.0000	0.0972	0.1473
Household owns one or more dining sets	0.2196	0.0221	9.9200	0.0000	0.1762	0.2629
Household owns one or more wardrobes or cabinets	0.1529	0.0163	9.3600	0.0000	0.1209	0.1849
Household has consumed meat in the past 7 days	0.0246	0.0020	12.1800	0.0000	0.0206	0.0286

Annex 2: Out-of-Sample Accuracy Tests

Note: The out-of-sample accuracy tests were not re-run using the new $1.25/day poverty line. However, we would expect the tool to exhibit even greater accuracy when predicting from a population with a higher rate of extreme poverty.

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 Cambodia tool was divided randomly into two data sets: the analysis data set with 11,238 households and a holdout data set with 3,746 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 very well, actually gaining in terms of BPAC. On PIE, the model performs almost identically well out-of-sample.

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

	Total Accuracy	Poverty Accuracy	Under-coverage	Leakage	PIE	BPAC
In-Sample Prediction						
	84.02	48.25	51.75	53.58	0.28	46.42
Out-of-Sample Prediction						
	84.50	49.12	50.88	49.03	-0.29	47.26

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

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

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 lower bound is defined by the 2.5th percentile of the sample distribution for each accuracy measure; the upper bound is defined by the 97.5th percentile (note that this method does not assume normality).

Table 8: Bootstrapped Confidence Intervals Computed Empirically from Sampling Distribution Without Normality Assumption

Accuracy Measure	95% Confidence Interval	
	LB	UB
Total Accuracy	83.36	85.69
Poverty Accuracy	44.84	53.45
Undercoverage	46.55	55.16
Leakage	41.90	56.38
PIE	-1.61	1.02
BPAC	36.80	51.51

As with the naïve result in Table 7, the performance of this model is quite good. BPAC has a relatively narrow confidence interval. For PIE, which measures the difference between the actual poverty rate in the sample and the predicted poverty rate, the confidence interval is especially tight. Taken together, the results from Tables 7 and 8 allow us to conclude that the 1-step Quantile model on which the Cambodia PAT is based performs well out-of-sample.