IDS Working Paper 161

Poverty monitoring and targeting using ROC curves: examples from Vietnam

Bob Baulch*

August 2002

INSTITUTE OF DEVELOPMENT STUDIES
Brighton, Sussex BN1 9RE
ENGLAND

* Fellow, Institute of Development Studies at the University of Sussex.

ISBN 1 85864 428 3

Summary

This paper suggests a low cost methodology for identifying and assessing the accuracy of poverty monitoring and targeting indicators. It proposes that the non-parametric technique of Receiver Operating Characteristic (ROC) curves be used to assess the accuracy of poverty indicators. Furthermore, it suggests that the best individual indicators can be combined into a composite poverty indicator using a stepwise Probit approach. The composite poverty monitoring indicators so developed are intuitive, parsimonious (involving just six to nine indicators in the examples from Vietnam), and relatively easy to collect data on. A further advantage of the method is that it allows the trade-off between coverage of the poor and exclusion of the non-poor to be quantified in terms that are readily understandable by policy-makers. This methodology (suitably adapted and expanded) could help bridge the gap between the aggregate poverty statistics generated by periodic household surveys and the need for more disaggregated and regular poverty statistics by welfare agencies.

Acknowledgements

The author thanks Stephen Devereux , Henry Lucas and Barry Reilly for helpful comments on an earlier version of this paper and Jenny Edwards for excellent secretarial assistance. Special thanks are due to Nick Minot for early methodological discussions and suggestions on ROC curves and stepwise Probits. The usual disclaimers apply.

Contents

Figures

Tables

1 Background and policy context

In most developing countries, it is only feasible to conduct detailed household sample surveys every three to five years. Yet the needs of government and donor agencies often require that poverty is monitored on an annual (or even intra-annual!) basis. Meanwhile, welfare agencies often need to devise equitable and transparent methods for identifying which individuals or households should qualify for anti-poverty programmes. There is no particular reason why these two tasks (monitoring of poverty and targeting the poor) need be linked, but in practice they often are.

The Socialist Republic of Vietnam illustrates the demands for regular poverty data well. Since the *doi moi* (renovation) reforms of the late 1980s, two high-quality household surveys (the Vietnam Living Standards Surveys) have been conducted by the General Statistics Office (GSO). Yet the Government's ten-year Socio-Economic Strategy and five year Development Plan, specify annual targets for the reduction in the percentage of households below the (food) poverty line. The GSO has responded to this need by fielding abbreviated Multi-Purpose Household Surveys in the years in which Living Standard Surveys have not been conducted. But the preliminary results of these surveys have usually not been available until more than a year after fieldwork, and are regarded as of dubious quality by many donor agencies (Glewwe and Yansaneh 2000). Meanwhile, the main Vietnamese welfare agency, the Ministry of Labour, Invalids and Social Assistance (MOLISA) conducts an annual exercise in which its district and commune representatives together with Party cadres makes lists of the poor households in each of the 10,000+ communes in the country. These lists are used for: (a) identifying which households qualify for targeted assistance (free or subsidised schooling, health insurance cards, exemptions from various commune taxes) at the local level, and (b) monitoring the effectiveness of the Hunger Eradication and Poverty Reduction (HEPR) Program at the provincial and national levels. Adherence to MOLISA procedures for identifying poor households at the commune level are known, however, to be highly variable (Conway 2001) and complaints about the accuracy and transparency of the lists are common. In the light of these problems, the IMF and World Bank's Joint Staff Assessment of the recently approved Vietnam Interim Poverty Reduction Strategy Paper stated that 'A full PRSP will require a significant amount of work in identifying indicators and establishing mechanisms to monitor them' (IMF/World Bank 2001). In short, there is an almost insatiable demand for poverty-related statistics by Government and donors that cannot be met by either the GSO's regular programme of household sample surveys or the MOLISA reporting system.

This paper suggests a low-cost methodology for bridging the gap between the poverty statistics generated by periodic, high-quality household surveys and the need for more frequent and more disaggregated poverty monitoring and targeting indicators. It suggests that the non-parametric technique of Receiver Operating Characteristic (ROC) curves be used to assess the accuracy of individual and composite poverty indicators. Furthermore, it is suggested that the most accurate poverty indicators, which number just six to nine variables in the case of Vietnam, can be used both to monitor poverty trends in the years between household sample surveys and to identify households worthy of receiving targeted assistance.

2 Approaches to poverty monitoring and targeting in other countries

As mentioned above, a low-cost, easy-to-use and transparent method for identifying the poor is needed by both the agencies responsible for poverty monitoring and the agencies responsible for the targeting of anti-poverty programmes. A brief survey of the developing country literature reveals a number of solutions to this need. In sub-Saharan Africa, a number of countries have also been working with the World Bank economists and statisticians to develop a Core Welfare Indicators Questionnaire (CWIQ) for monitoring the effect of policies, programmes and projects on poverty and living standards. The 'CWIQ is intended to be applied frequently (possibly annually)' and 'to provide rapid information on key indicators for policy target groups' (World Bank 1999: 2 and 7). Other countries, such as Ghana and Uganda, have used regression-based techniques to identify indicators for monitoring poverty over time based on their own periodic household sample surveys (Fofack 2000; MFPED 2001). Donors and NGOs have also long-supported the development of food security information systems in Sub-Saharan Africa (Devereux 2001). Some well-known examples are, FAO's Food Insecurity and Vulnerability Information and Mapping Systems (FIVIMS), Save the Children UK's Food Economy/RiskMap approach, USAID's Famine Early Warning System and the World Food Programme's Vulnerability Assessment and Mapping reports.

In Latin America, there is a long tradition of targeting social programmes using "proxy means tests" (Grosh and Baker 1995). In Chile, for example, social workers in many municipalities are given simple forms to collect information on 14 household characteristics, which are used to determine whether a household qualifies for welfare benefits (such as disability and old age pensions, family subsidy and housing benefits). Simulations based on household survey data from Bolivia, Jamaica, and Peru show that 'household characteristics can serve as reasonable proxies for income in assessing eligibility for social programmes' although 'significant errors of under-coverage' invariably occur (Grosh and Baker 1995: ix). In a related paper, based on data from 38 social sector programs in 11 Latin America countries, Grosh finds that 'targeted programs have a much more progressive incidence than general food price subsidies' (Grosh 1995: 483). Furthermore, 'the administrative costs of programs with moderately good incidence need not be excessively high' (Grosh *op. cit.*).

In South Asia, there is more of a tradition of self-targeting anti-poverty programmes such as the well-known Maharastra Employment Guarantee Scheme in India, food stamps in Sri Lanka and "food for education" in Bangladesh. Nonetheless, work has been done in Bangladesh (Rahman and Hossain 1995; Wodon 1997) and Pakistan on developing poverty and well-being indicators using household sample surveys. In South India, Ravallion and Chao (1989) have applied a quadratic programming approach to household level panel data collected by ICRISAT to determine an optimal distribution of targeting benefits.[1] As discussed below, this approach while theoretically sound has the disadvantage of being complex and expensive to implement on a regular basis.

3 Evaluating the accuracy of individual poverty indicators

We assess the accuracy of different poverty indicators using a relatively novel technique: Receiver Operating Characteristic (ROC) curves, a graphic and non-parametric way of portraying the ability of different diagnostics tests to distinguish between a binary outcome originally developed for use in electrical engineering and signal-processing (Stata Corp 2001).[2] A ROC curve shows the ability of a diagnostic test to distinguish correctly between two states or conditions. In the context of poverty targeting, a ROC curve plots the probability of a test identifying a person as poor (known as the test's "sensitivity") on the vertical axis against 1 minus the probability of the same test correctly classifying a person as non-poor on the horizontal access (known as the test's "specificity"). When the diagnostic test (here the value of the poverty indicator) can take several discrete values, its ROC curve will consist of a series of linear segments corresponding to these discrete values. The greater the area under an ROC curve, which looks like an inverted Lorenz curve, the greater is the efficacy of a diagnostic test. Conversely, the closer a ROC curve is to the 45 degree line, the weaker is its efficacy.

Figure 3.1 shows an example of a ROC curve drawn from household survey data from Vietnam. The six segments of the curve correspond to the six different type of floor observed in a nationally representative household survey, the Vietnam Living Standards Survey of 1997–98 (hereafter VLSS98). The vertical axis shows the extent to which different floor types allow one to classify poor people correctly as poor (the test's sensitivity) using an expenditure poverty line of VND 1,789,871 per person per year.[3] The horizontal axis, when read from right to left, shows the extent to which different floor types, which have been ordered by their likely association with poverty, allow non-poor people to be correctly identified (the specificity of a test). However, in order to show the usual trade-off between coverage of the poor and leakages to the non-poor, the usual format for a ROC curve is to plot sensivity against 1- specificity.

Consider the first and lowest segment of the curve, which corresponds to people living in houses with earth floors (some 32 per cent of the total population). If all people living in houses with this simplest type of flooring were classified as poor, then just over half (51 per cent) of poor people would be identified. However, over one-fifth (22 per cent) of non-poor people also live in households with earth floors. Now consider the second segment of the ROC curve, which corresponds to wooden floors. If all people living in houses with earth and wooden floors (some 38 per cent of the population) were

---

1 Glewwe (1992) takes a similar optimization approach, though using non-linear mathematical programming, to poverty targeting in Cote D'Ivoire.

2 To our knowledge, the only previous use of ROC for analysing the impact of poverty targeting was by Wodon (1997) using household survey data from Bangladesh.

3 This "overall poverty" line was established by the Vietnamese General Statistics Office in consultation with the World Bank and is based on VLSS98. See footnote 10 for further details.

considered poor, the percentage of the poor covered would increase to about three-fifths (59 per cent) but at the expense of around a quarter of non-poor people also being classified as poor. As higher quality floor types (made, respectively, of lime, cement, bricks or tiles) are successively included, so the coverage of the poor increases but at the expense of more and more non-poor people being wrongly included.

Figure 3.1 ROC curve for floor types in Vietnam

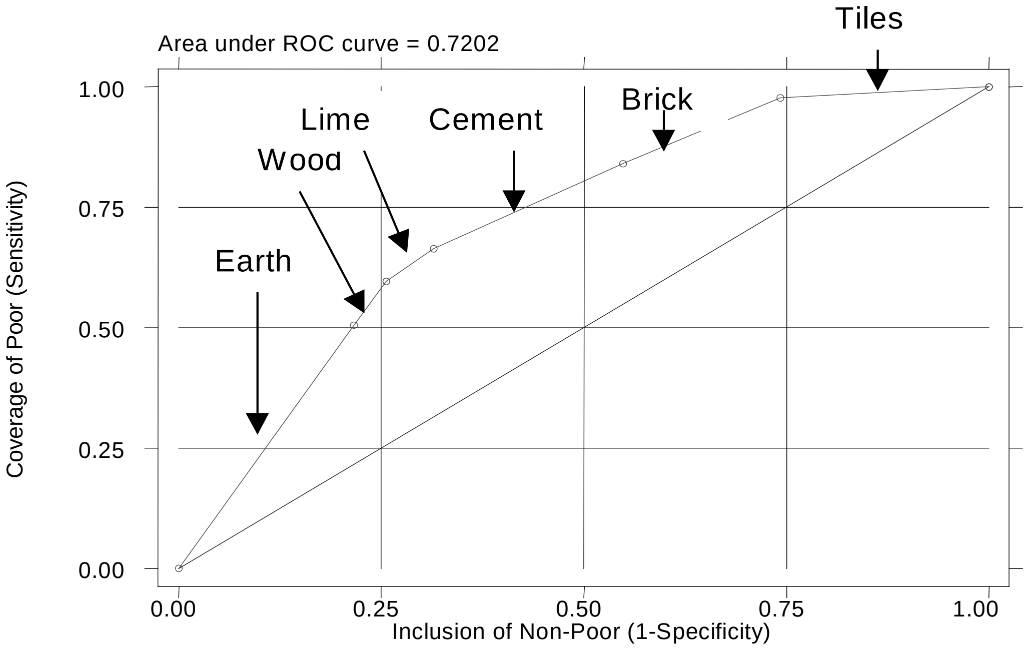

Several additional points can be made using this illustration. First, as shown in Table 3.1 below, choosing the categories (earth+wood) which correspond to the highest percentage of correctly classified poor and non-poor people is not ambiguously the best cut-off. Some policy-makers might argue that it was better to err on the side of caution and also include those living in houses with lime floors as poor, in which case two-thirds of poor people would be correctly classified. On the other hand, expenditure "hawks" who were keen to exclude as many people as possible from programme benefits might argue for only those living in houses with earth floors, in which case leakages to the non-poor would be minimised. Clearly, this reflects a welfare judgement which economists are poorly equipped to make. However, ROC curves (and their accompanying tables) provide a useful way of summarising the trade-off.

Second, ROC curves can be linked to the occurrence of Type I and Type II errors familiar from conventional statistical hypothesis testing (known as "false positives" and "false negatives" in epidemiology and medicine) as follows.[4] Sensitivity is 1 minus the probability of a Type I error (incorrectly classifying a poor households as non-poor) while 1 minus the specificity of a test is the same as the probability of a Type II error (incorrectly classifying a non-poor household as poor). In many respects

[4] Following Cornia and Stewart (1995), Type II errors are also known as E errors (as they involve "excessive program coverage") and Type II errors as F errors (as they involve a "failure" to reach the target population).

these differences in terminology are akin to describing whether "a glass is half-empty or half-full", in that both are simply different methods of analysing the same phenomenon.

Table 3.1 Trade off between coverage of the poor and inclusion of the non-poor

Type of floor	Coverage of poor (Sensitivity)	Inclusion of non-poor (1-Specificity)	% Correctly classified
Earth	50.5%	21.7%	67.9%
Earth + wood	59.6%	25.6%	68.8%
Earth + wood + lime	66.4%	31.4%	67.8%
Earth + wood + lime + cement	84.1%	54.8%	59.7%
Earth + wood + lime + cement + brick	97.7%	74.3%	52.6%

Source: Author's calculations from VLSS98.

Third, as long as a potential poverty indicator is "monotonically increasing with the risk of failure" (i.e., it increases in value as the likelihood of poverty increases), then the area under a ROC curve can be used for ranking the efficacy of different poverty indicators (Stata Corp 2001). The more a test's ROC curve is bowed toward the upper left-hand corner of the graph, the greater is the accuracy of the test. Since the ROC curves are bounded by the interval [0,1], the maximum value for the area under an ROC curve is 1 (in which case the test would predict poverty perfectly and the ROC curve would coincide with the left-hand vertical and top horizontal axes).[5]

Finally, it should be noted that ROC analysis can only be employed for dichotomous outcome variables (so that it can be used for the conventional headcount index of poverty but not for higher-order measures of poverty such as the poverty-gap and squared poverty gap).[6]

4 Methodology for producing a composite poverty indicator

It will usually be the case that some combination of indicators will provide a better poverty monitoring indicator than any single indicator. A number of previous studies have proposed different methodologies for producing such a composite indicator (usually in the context of targeting an anti-poverty alleviation). These include using regression analysis to predict per capita expenditures (Grosh and Baker 1995; MFPED 2001), linear and quadratic programming (Glewwe 1992; Ravallion and Chao 1989) and principal components (Zeller *et al.* 2001). Linear regression analysis, often performed after transforming expenditures or incomes into logarithmic terms, is relatively easy to implement but has the disadvantage that the objective function minimised relates to the entire income or expenditure distribution rather than

5 In contrast, a test with no predictive power would correspond to an area of 0.5 under the ROC curve (which would itself coincide with the 45 degree line in the ROC diagram).

6 The poverty gap and squared poverty gap are commonly used quantitative measures of, respectively, the depth and intensity/severity of poverty. See Ravallion (1994) for further details.

the distribution among the poor. Linear or quadratic programming avoids the problem but is more complex to implement, is less transparent to policy-makers, and requires that the monetary amount available for anti-poverty interventions is specified in advance. Principal components and factor analysis are efficient methods for combining different indicators, but there is no reason why the principal components and factor scores produced should be a good proxy for poverty.

We propose an alternative method for identifying a composite, and yet parsimonious, poverty indicator: estimating a stepwise Probit and then assessing its ability to distinguish the poor from the non-poor using ROC analysis.[7] As is well known, Probits are used when the dependent variable is dichotomous or polychotomous, while stepwise regression adds or deletes individual variables from the estimating equation in an iterative fashion according to their significance levels ("t" or "z" statistics) until certain pre-specified criteria is satisfied. Combining a Probit with backwards stepwise selection is therefore a parsimonious way of letting the data determine which household and other characteristics are good (joint) predictors of poverty. Furthermore, as explained below, by varying the cut-off for the probability of poverty, it is possible to determine the extent to which the poverty indicator covers the poor at the expense of including the non-poor. ROC analysis provides a method for displaying and assessing this trade-off. With recent advances in computing software, it is now very easy to perform this kind of analysis.[8]

Before describing how this methodology is implemented, it is important to stress that a sensible composite poverty indicator cannot be developed by blind adherence to econometric and statistical techniques. As with any econometric or statistical model, the policy analyst needs to exercise a good deal of judgement and common sense in estimating and evaluating the usefulness of different poverty indicators. Like programming and principal components, the stepwise Probit is designed to find the set of indicators that satisfy all the pre-determined selection (or de-selection) criteria. Whether or not the set of poverty indicators that are identified make "economic sense", whether they are sufficiently parsimonious, and if these indicators are relatively easy to collect reliable information on, must be left up to the analyst's judgement.[9]

---

7 Discriminant analysis or classification methods might also be used to derive discriminant functions or classification statistics for distinguishing the poor from the non-poor, without many of the distributional assumptions required by a stepwise Probit.

8 The stepwise Probit estimation and ROC analysis contained in this paper were conducted using Stata, version 7.0.

9 For example, indicators which can change rapidly over time (such as the ownership of consumer durables or vehicles) can work well for monitoring poverty but may be problematic targeting indicators because it is relatively easy for potential welfare recipients to conceal ownership of these items.

5 Empirical application to Vietnam

5.1 Individual indicators

Most household surveys contain a large number of possible poverty indicators at the household level together with other community or geographic variables that may be good predictors of poverty. The Vietnamese Living Standards Surveys (VLSS) of 1997–98, standard household sample surveys patterned after the World Bank's Living Standards Measurement Survey, are no exception to this with data collected on over 3,000 variables.[10] We here examine the area under the ROC curve for a sub-set of these variables, most of which were identified as significant correlates of poverty in a separate poverty mapping exercise that combined data from the 1999 Census and 1998 VLSS (Minot and Baulch 2002). Due to the substantial differences between rural and urban areas, we conducted this analysis for Vietnam as a whole and for rural communes and urban wards considered separately. Two poverty lines, corresponding to the minimum expenditure needed to acquire 2100 Kcals with or without a modest allowance for non-food expenditures, were also used.[11]

Table 5.1 shows the area under the ROC curves for five groups of individual poverty indicators for rural communes, urban wards and all of Vietnam. All of these poverty indicators are categorical variables, on which it would be relatively easy to collect reliable information in a poverty monitoring survey.[12] Recall, following on from the discussion in Section 3, that the closer the area under the ROC curve is to 1, the greater is its efficacy as a poverty indicator.

Household characteristics generally perform poorly as poverty indicators in rural areas, although the number of children under 15 years of age (the age by which Vietnamese children should have completed lower secondary school) in a household is a good indicator of food poverty. In urban areas, the educational levels completed by the household head and spouse (which are themselves highly correlated) are also good indicators of poverty along with the number of children in a household. The ethnicity of the household head is a reasonable indicator of both food and overall poverty in rural areas, but performs poorly in urban areas where very few ethnic minority people live. For Vietnam as a whole, the number of children under 15 emerges as the best poverty indicator followed by the number of females and ethnicity of the household head.

Moving on to housing quality, floor type is a uniformly better indicator of both food and overall poverty than roof type. In urban areas, toilet type is also a good poverty indicator where the main cooking fuel and source of drinking water are also extremely good poverty indicators. The areas under the ROC

---

10 See Haughton (2000) for an excellent description of the VLSS.

11 The first of these poverty lines, the "food poverty line", corresponds to the expenditure necessary to obtain 2,100 Kcals per person per day (VND 1,286, 8333). The second "overall" poverty line corresponds to the food poverty line plus a modest allowance for essential non-food expenditures (VND 1,789,971). In 1998, US$1 ≙ VND 14,000. See Appendix 2 of Poverty Working Group (1999) for further details on the calculation of these poverty lines.

12 A further advantage of categorical indicators is that they are less susceptible to measurement error than the continuous variables (such as expenditures, incomes or land holdings) on which poverty measurement is usually based.

Table 5.1 Accuracy of different indicators in targeting the poor in Vietnam

Indicator	Area under Receiver Operating Characteristic Curve					
	Rural		Urban		All Vietnam	
	Food Poverty	Overall Poverty	Food Poverty	Overall Poverty	Food Poverty	Overall Poverty
Household Characteristics						
Educational Level of Household Head	0.601	0.579	0.715	0.685	0.625	0.609
Educational Level of Spouse *	0.570	0.554	0.739	0.727	0.602	0.597
Number of Children under 15	0.733	0.690	0.753	0.789	0.742	0.714
Number of Females	0.636	0.618	0.578	0.671	0.632	0.616
Ethnicity of Household Head	0.642	0.612	0.495	0.500	0.649	0.614
Housing Quality						
Floor Type	0.696	0.665	0.694	0.773	0.734	0.720
Roof Type	0.630	0.585	0.687	0.658	0.637	0.594
Toilet Type	0.597	0.577	0.773	0.730	0.650	0.648
Fuel and Water						
Cooking Fuel	0.585	0.570	0.759	0.795	0.641	0.650
Source of Drinking Water	0.580	0.577	0.765	0.730	0.641	0.652
Durable Assets						
Irrigated Land Allocated	0.529	0.542	n/a	n/a	0.619	0.646
Radio and TV Ownership	0.736	0.711	0.876	0.792	0.771	0.751
Vehicle Ownership	0.649	0.643	0.763	0.773	0.758	0.677
Geographic						
Poor or Remote Commune	0.585	0.559	0.554	0.520	0.589	0.559
Geograpic Region	0.666	0.622	n/a	n/a	0.726	0.707

Notes on Indicators:

Educational Level Completed (both Heads and Spouses): 0 = Post-secondary; 1=Advanced Technical; 2=Upper Secondary; 3=Lower Secondary; 4=Lower Secondary; 5=Primary; 6=Less than Primary (* Note: 1284 households do not have spouses present)

Ethnicity: 0=Kinh or Chinese Head; 1= Ethnic minority head

Floor Type: 0=Earth; 1=Other, 2=Bamboo/Wood; 3=Lime and Ash; 4=Cement; 5=Brick; 6=Marble or Tile

Roof Type: 0=Other; 1=Leaves/Straw; 2=Bamboo/Wood; 3=Canvas/Tar Paper; 4=Panels; 5=Galvanised Iron; 6=Tile; 7=Cement or Concrete

Toilet Type: 0=Flush; 1=Latrine/Other; 2=No toilet

Cooking Fuel: 1=Bottled Gas; 2=Electricity; 3=Kerosene; 4=Coal/Charcoal; 5=Leaves/Straw; 6=Wood

Source of Drinking Water: 1=Tap; 2=Deep Well with Pump; 3=Other Well or Spring; 4=Rainwater; 5=Lake or River; 6=Other

Irrigated Land: Quintiles of irrigated land allocated (Note: 1 is top quintile and 5 is bottom quintile)

Radio and TV Ownership: 1= Color TV; 2=Radio; 3=Black & White TV; 4=Color TV

Vehicle Ownership: 1=Motorbike (or car); 2=Boat; 3=Bicycle; 4=No Vehicle

Poor commune: 0=Commune not included in CEMMA's list of difficult mountaineous and remote communes or MOLISA list of poor communes; 1=Commune included in either CEMMA difficult mountainous and remote communes or MOLISA poor communes lists;

Geographic Region: 1=Urban; 2=Midland; 3=Inland Delta; 4=Coastal; 5=Low Mountains; 6=High Mountains curves for cooking fuel and drinking water are, however, statistically indistinguishable from one another (using asymptotic normal confidence intervals).

Ownership of durable assets such as land, radios, televisions and vehicles are usually inversely related to poverty. In the case of Vietnam, the use of land as a poverty indicator is complicated by the fact that, by law, land is not owned but allocated to households for a fixed period (usually between 30 and 50 years) by the commune authorities. In addition, land quality varies widely between the high productive deltas and coastal and mountainous areas. As a consequence, irrigated land allocated performs rather poorly as a poverty indicator in rural communes. In contrast, ownership of radios, televisions and vehicles perform rather well as poverty indicators, especially in urban wards. This is a consequence of the widespread ownership of bicycles and inexpensive battery-operated black and white televisions (often imported from China) in rural areas, and the growing popularity of colour televisions and motorcycles in Vietnam's rapidly expanding towns and cities.

Finally, two geographic poverty indicators are considered. It can be seen that the Government of Vietnam's current system of "poor and remote communes" is not particularly accurate as a poverty indicator.[13] The simple reason for this is that the vast majority of poor people in Vietnam do not live in an officially designated poor or remote commune. The five types of geographic region identified in the VLSS do somewhat better at identifying poverty in rural areas and even better, once urban areas are added as a sixth category, for Vietnam as a whole. It is not, however, possible to use geographic region as an indicator of urban poverty, as all major cities in Vietnam are located in inland deltas or coastal areas.

Poverty is, of course, a multidimensional concept. Thus while it is operationally convenient to define the poor as those with per capita expenditures less than the poverty line, it is also useful to examine other wider definitions of ill-being. Table 5.2 shows the area under the ROC curves for the same indicators as in Table 5.1 but with malnutrition and illiteracy of the household head /spouse used to define poverty/ill-being. While the ranking of the accuracy of the indicator variables are broadly preserved, the area under the ROC curves for these wider definitions of ill-being are substantially lower than for the poverty line approach. This illustrates Sen's (1985) thesis that, because of differences in capabilities, standard measures of material poverty (such as the headcount) are imprecise measures of poverty outcomes.

[13] Vietnam has two official lists for identifying poor and remote communes. Under Programme 133, the Ministry of Labour, Invalids and Social Assistance maintains a list of "poor communes", most of which are located in lowland areas. Under Programme 135, a separate list of "especially difficult mountainous and remote communes" is kept by the Committee for Ethnic Minorities in Mountainous Areas. These lists have been combined to create the poor and remote communes indicator used in this paper.

Table 5.2 Accuracy of different variables in identifying 'ill-being' in Vietnam

Indicator	Area under Receiver Operating Characteristics Curve							
	Malnourished Head (BMI < 18.5)		Malnourished Spouse		Illiteracy of Head		Illiteracy of Spouse	
	Rural	Urban	Rural	Urban	Rural	Urban	Rural	Urban
Household Characteristics								
Educational Level of Household Head	0.535	0.561	0.508	0.597	0.867	0.918	0.729	0.7997
Educational Level of Spouse	0.539	0.606	0.506	0.579	0.729	0.800	0.854	0.9153
Ethnicity	0.480	0.501	0.500	0.506	0.588	0.509	0.569	0.5022
Housing Quality								
Floor Type	0.514	0.560	0.525	0.610	0.626	0.559	0.609	0.5435
Roof Type	0.497	0.513	0.481	0.533	0.658	0.597	0.638	0.5987
Toilet Type	0.498	0.544	0.500	0.584	0.625	0.547	0.591	0.5331
Fuel and Water								
Cooking Fuel	0.492	0.596	0.481	0.632	0.569	0.569	0.559	0.5421
Source of Drinking Water	0.501	0.563	0.493	0.570	0.572	0.534	0.556	0.5069
Durable Assets								
Irrigated Land Allocated	0.507	n/a	0.496	n/a	0.576	n/a	0.573	n/a
Radio and TV Ownership	0.552	0.583	0.551	0.571	0.630	0.612	0.597	0.6075
Vehicle Ownership	0.552	0.569	0.542	0.608	0.630	0.661	0.615	0.6426
Geographic								
Poor or Remote Commune	0.490	0.523	0.513	0.522	0.552	0.492	0.538	0.4899
Geographic Region	0.506	n/a	0.524	n/a	0.521	n/a	0.509	n/a

Notes on Indicators:

Educational Level Completed (both Heads and Spouses): 0 = Post-secondary; 1=Advanced Technical; 2=Upper Secondary; 3=Lower Secondary; 4=Lower Secondary; 5=Primary; 6=Less than Primary (* Note: 1284 households do not have spouses present)
Ethnicity: 0=Kinh or Chinese Head; 1= Ethnic minority head
Floor Type: 0=Earth; 1=Other, 2=Bamboo/Wood; 3=Lime and Ash; 4=Cement; 5=Brick; 6=Marble or Tile
Roof Type: 0=Other; 1=Leaves/Straw; 2=Bamboo/Wood; 3=Canvas/Tar Paper; 4=Panels; 5=Galvanised Iron; 6=Tile; 7=Cement or Concrete
Cooking Fuel: 1=Bottled Gas; 2=Electricity; 3=Kerosene; 4=Coal/Charcoal; 5=Leaves/Straw; 6=Wood
Source of Drinking Water: 1=Tap; 2=Deep Well with Pump; 3=Other Well or Spring; 4=Rainwater; 5=Lake or River; 6=Other
Toilet Type: 0=Flush; 1=Latrine/Other; 2=No toilet
Irrigated Land: Quintiles of irrigated land allocated (Note: 1 is top quintile and 5 is bottom quintile)
Radio and TV Ownership: 1= Color TV; 2=Radio; 3=Black & White TV; 4=Color TV
Vehicle Ownership: 1=Motorbike (or car); 2=Boat; 3=Bicycle; 4=No Vehicle
Poor commune: 0=Commune not included in CEMMA's list of difficult mountaineous and remote communes or MOLISA list of poor communes; 1=Commune included in either CEMMA difficult mountainous and remote communes or MOLISA poor communes lists;
Geographic Region: 1=Urban; 2=Midland; 3=Inland Delta; 4=Coastal; 5=Low Mountains; 6=High Mountains

5.2 Composite indicators

The stepwise Probit methodology described above was used to construct composite poverty indicators for rural communes and urban wards. Based on the analysis of individual poverty indicators in Table 5.1, one might expect that the housing quality, ownership of durable assets, and some household characteristic (such as education and the number of children or women in the household) would be good variables to include in the composite poverty indicators. However, some of these individual indicators (such as floor and roof type, or ethnicity and geographic region) might be so highly co-linear with each other that collectively they add little to an indicator's predictive accuracy. Estimation of stepwise Probits (with backward selection of variables) allows the contribution of each variable to be sequentially evaluated and compared, and a parsimonious list of variables for the computation of the composite poverty indicator identified. Care must, however, be taken to avoid over-fitting of the data, so that a few variables selected by the stepwise regression may need to be excluded on *a priori* grounds by the analyst. It is also sometimes necessary to merge some categories of indicator variables together to avoid over-determination problems. Appendix 1 details the procedure used to derive a composite poverty indicator for rural communes, whose final estimating equation is summarised below:

Rural communes

Poverty Indicator = **Φ** (0.27 x number of children in household
+ 0.17 x number of women in household
+ 0.74 if household head comes from an ethnic minority
- 0.87 if there is a colour TV in household
- 0.44 if there is a black-and-white TV in household
- 0.32 if household owns a radio
- 1.35 if someone in household owns a car/motorbike
+ 0.20 if house floor is made of earth
+ 0.63 if main cooking fuel is leaves, straw or wood
- 1.35)

where **Φ** is the cumulative normal distribution. The composite poverty indicator variable this equation produces will vary between 0 and +1 and corresponds to the probability that an individual is poor. The area under the ROC curve corresponding to the composite poverty indicator for rural communes is 0.847.

The corresponding composite poverty indicator for urban wards (in which ethnic and geographic terrain are not included in the initial list of variables fed into the stepwise Probit) is as follows:

Urban wards

Poverty Indicator = Φ (0.34 x number of children in household
+ 0.19 x number of females in household
+ 1.10 if main cooking fuel is charcoal, coal, leaves, straw or wood
- 0.53 if there is a black and white TV in household
- 0.96 if there is a colour TV in household
- 1.32 if someone in households owns a car/motorbike
- 1.65)

The area under the ROC curve for the composite poverty indicator for urban wards is 0.930.

Figures 5.1 and 5.2 show the ROC diagrams for rural communes and urban wards using the two composite poverty indicators developed above. Note how the ROC curve for the composite poverty indicator in urban areas rises much more sharply than the ROC curve for rural areas. This indicates the greater predictive accuracy of the urban poverty indicator despite the fact it is based on three less variables than the rural indicator.

Figure 5.1 Rural communes

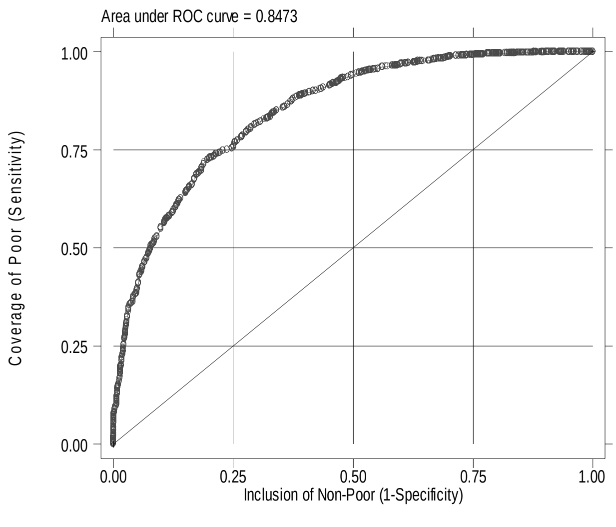

Figure 5.2 Urban wards

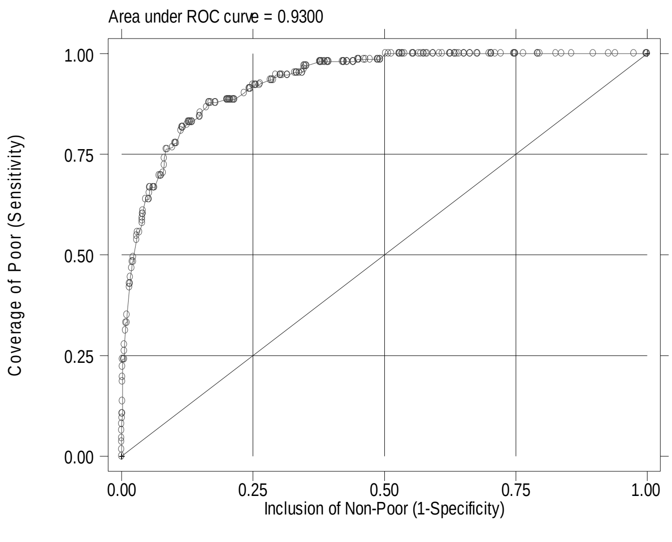

Having estimated the probability of poverty, the question then arises as to what cut-off point should be used to determine whether or not a person or household qualifies for government assistance? In selecting such a cut-off point there will always be a trade-off between coverage of the poor and exclusion of the non-poor. In general, the more accurately one is able to identify the poor, the less accurately one is able to identify the non-poor, and vice-versa.

Tables 5.3 and 5.4 show the trade-offs between correct coverage of the poor and exclusion of the non-poor in rural communes and urban wards. Consider Table 5.3, which shows the possible cut-offs of the composite poverty indicator for rural communes shown above. If a very low value for the cut-off such as 0.025 (row 1 in Table 5.3) were chosen, nearly all households (99.9 per cent) would be correctly identified as poor. In contrast if a very high value was chosen for the cut-off such as 0.975 (the last row of Table 5.3), nearly all households (99.9 per cent) would be correctly identified as non-poor. Clearly, a cut-off for the probability of poverty needs to be selected which is between these two values. The value chosen will depend on the relative importance that the policy-maker attaches to: (i) coverage of poor, and (ii) exclusion of the non-poor. In ROC terms, these two objectives correspond to sensitivity and specificity respectively.

When policy-makers attach equal weight to these two objectives, then the appropriate cut-off for the probability of poverty would be that which maximizes the percentage of correctly identified poor and non-poor households, as shown in column 4 of Table 5.3. This maximum occurs when the probability of poverty is 0.5. At this cut-off, almost three-quarters (76.5 per cent) of households are correctly identified as poor or non-poor.

If, however, policy-makers care more about reducing the number of poor people excluded from program benefits than about inclusion/leakages to the non-poor, then a lower cut-off for the probability of poverty might be chosen. For example, a cut-off of 0.325 would ensure that 89.8 per cent of the rural poor were correctly identified, although at the cost of excluding only 58.8 per cent of non-poor people. Hence, the overall percentage of people correctly identified as poor and non-poor falls to 72.7 per cent.

Conversely, if policy-makers attach greatest priority to avoiding leakages of poverty alleviation funds to non-poor households, they might pick a higher probability cut-off which increases the exclusion of the non-poor. For example, choosing a cut-off of 0.65 ensures that 89.9 per cent of the non-poor are excluded from program benefits. This increase in targeting accuracy comes, however, at the cost of correctly identifying just 55.3 per cent of poor people. Again, the overall percentage of households correctly identified falls (to 74.4 per cent).

Table 5.4 shows the trade-off between correctly identifying poor and non-poor households in urban wards. This table is interpreted in a similar way to Table 5.3, although – due to the steeper slope of the ROC curve for the urban composite poverty indicator – the trade-off between coverage of the poor and exclusion of the non-poor is less stark.

Table 5.3 Trade-off between coverage of the poor and exclusion of the non-poor in rural areas

Cut-off >=	Coverage of poor (Sensitivity)	Exclusion of non-poor (Specificity)	% Correctly classified
0.025	99.9%	10.4%	50.6%
0.050	99.6%	17.4%	54.3%
0.075	99.4%	22.9%	57.2%
0.100	99.1%	26.8%	59.2%
0.125	98.4%	30.9%	61.1%
0.150	97.7%	35.4%	63.3%
0.175	97.0%	38.8%	64.9%
0.200	96.5%	41.1%	65.9%
0.225	95.4%	45.7%	68.0%
0.250	94.7%	48.5%	69.2%
0.275	93.0%	52.1%	70.5%
0.300	91.7%	54.5%	71.2%
0.325	89.8%	58.8%	72.7%
0.350	88.0%	61.6%	73.5%
0.375	85.8%	64.1%	73.8%
0.400	83.4%	67.4%	74.6%
0.425	80.9%	71.0%	75.4%
0.450	77.7%	73.6%	75.4%
0.475	75.4%	76.8%	76.2%
0.500	73.1%	79.3%	76.5%
0.525	70.2%	81.4%	76.4%
0.550	68.0%	82.9%	76.2%
0.575	64.9%	84.6%	75.8%
0.600	60.7%	86.9%	75.1%
0.625	58.4%	88.3%	74.9%
0.650	55.3%	89.9%	74.4%
0.675	51.2%	91.5%	73.5%
0.700	48.5%	92.5%	72.8%
0.725	44.7%	93.8%	71.8%
0.750	41.1%	94.7%	70.7%
0.775	38.3%	95.5%	69.8%
0.800	34.1%	96.8%	68.7%
0.825	30.5%	97.2%	67.3%
0.850	26.6%	97.7%	65.8%
0.875	20.9%	98.3%	63.6%
0.900	17.6%	98.5%	62.2%
0.925	13.6%	99.3%	60.8%
0.950	9.0%	99.8%	59.1%
0.975	5.7%	99.9%	57.7%

Table 5.4 Trade-off between coverage of the poor and exclusion of the non-poor in urban areas

Cut-off >=	Coverage of poor (Sensitivity)	Exclusion of non-poor (Specificity)	% Correctly classified
0.025	97.9%	57.8%	61.4%
0.050	94.7%	69.8%	72.0%
0.070	91.3%	75.7%	77.1%
0.100	88.6%	80.0%	80.8%
0.125	85.3%	85.1%	85.1%
0.150	83.1%	87.0%	86.7%
0.174	80.9%	88.7%	88.0%
0.200	76.3%	91.3%	89.9%
0.224	70.3%	92.1%	90.1%
0.250	66.7%	93.7%	91.3%
0.275	66.7%	94.0%	91.5%
0.300	63.8%	94.8%	92.0%
0.324	61.0%	95.9%	92.8%
0.349	58.6%	96.1%	92.7%
0.375	55.6%	96.6%	92.9%
0.400	54.7%	97.1%	93.3%
0.425	53.7%	97.1%	93.2%
0.449	49.4%	97.8%	93.4%
0.474	48.3%	98.0%	93.5%
0.500	46.7%	98.1%	93.4%
0.524	44.4%	98.3%	93.5%
0.550	42.8%	98.3%	93.3%
0.574	35.1%	99.0%	93.2%
0.625	33.1%	99.2%	93.2%
0.649	27.7%	99.5%	93.0%
0.675	24.0%	99.5%	92.6%
0.699	22.2%	99.8%	92.8%
0.725	19.6%	99.8%	92.6%
0.750	18.5%	99.8%	92.4%
0.800	13.6%	99.8%	92.0%
0.850	10.6%	99.8%	91.7%
0.875	9.5%	99.9%	91.7%
0.899	6.5%	100.0%	91.5%
0.925	4.5%	100.0%	91.4%

6 Conclusions and caveats

This paper has suggested an easy-to-use and low-cost method for identifying poverty monitoring and targeting indicators. It uses the non-parametric technique of Receiver Operating Characteristic curves to access the accuracy of individual poverty indicators, and then combines the best individual indicators using a stepwise Probit approach. The composite poverty indicators developed are easy-to-compute, parsimonious (involving just six to nine indicators in the case of Vietnam), and easy to collect data on. A further advantage of this method is that it allows the trade-off between coverage of the poor and exclusion of the non-poor to be quantified in terms that are readily understandable by policy-makers.

A number of caveats should, however, be mentioned. A first caveat is that these indicators have been developed using a static cross-section only. If a large number of households move in and out of poverty over time (see Baulch and Hoddinott 2000) and policy-makers only want to target the chronically poor, then a method of identifying the chronically poor would be needed. In the Vietnam context, one possibility would be to use the panel sub-component of the VLSS to identify a poverty profile for the chronically poor along the lines of Gibson's (2001) study for Papua New Guinea.

A second caveat concerns the neglect of the administrative costs of targeting, which Grosh (1995) found varied between 6 and 9 per cent of total program costs in Latin America. In general, as Besley and Kanbur (1993) argue, it seems likely that the more finely an anti-poverty programme is targeted, the higher its administrative costs will be. Nonetheless, in a country as geographically diverse as Vietnam there may well be a case for disaggregating further and developing composite poverty indicators for individual regions or provinces.[14]

Third, the stepwise Probit approach used to select the individual indicator to use for producing a composite poverty indicator is one of a number of statistical methods which might be used. Discriminant analysis or classification methods, both which are well-established multivariate methods for analysing data from two (or more) populations, offer promise in this regard.[15]

A final caveat concerns the fact that a high level of inclusion/leakages to the non-poor may be necessary to secure sufficient political support for the anti-poverty intervention. As Gelbach and Pritchett (2000) put it: 'A leaky bucket may be better for redistribution to the poor'.

[14] The expanded sample size of the forthcoming 2002 Multi-Purpose Household Survey should allow this to be done.

[15] This is the subject of on-going research by the author.

Appendix 1: Estimation of composite poverty indicator for rural communes

This appendix describes the estimation procedure used for deriving the composite poverty indicator for rural communes using the 1997–98 Vietnam Living Standards Survey in Stata version 7.0. First, a stepwise Probit with analytic weights and backward selection of variables is used to identify the indicators that best predict poverty. 8 of 18 indicators (46 of 56 dummy variables) are removed. See the notes to Table 5.1 for how these variables are coded. The selection of the remaining ten variables (8 dummies and 2 continuous variables) is checked using the SvyProbit command (which takes account of the clustering and stratification of the survey design in the calculation of coefficient standard errors). This allows one indicator (source of drinking water) to be eliminated and two indicators (cooking fuel and floor type) to be refined. Finally, the Probit is re-estimated, a composite poverty indicator computed, and the area under its ROC curve calculated.

Note that a forward stepwise Probit identifies the same nine indicators plus land quintile 5 as the backward stepwise Probit. A high nominal level of statistical significance (P-value=0.001) is specified for the removal or inclusion of variables as a crude way to adjust for the pre-test bias that arises from the repeated testing of the same data.

```
.set matsize 400

. xi: sw probit allpoor morc i.georeg i.edchd2 i.edcsp2 children numfem ethnic i.floor2 i.roof2 toilet2 i.water2
i.fuel3 i.qtland radio bwtv colortv carmb bicycle [aw=hhsizewt] if urban92==0, pr(0.0001)
i.georeg        _Igeoreg_1-6        (naturally coded; _Igeoreg_1 omitted)
i.edchd2        _Iedchd2_0-5        (naturally coded; _Iedchd2_0 omitted)
i.edcsp2        _Iedcsp2_0-5        (naturally coded; _Iedcsp2_0 omitted)
i.floor2        _Ifloor2_1-6        (naturally coded; _Ifloor2_1 omitted)
i.roof2         _Iroof2_1-5         (naturally coded; _Iroof2_1 omitted)
i.water2        _Iwater2_1-6        (naturally coded; _Iwater2_1 omitted)
i.fuel3         _Ifuel3_3-6         (naturally coded; _Ifuel3_3 omitted)
i.qtland        _Iqtland_1-5        (naturally coded; _Iqtland_1 omitted)
                begin with full model
p = 0.9484 >= 0.0001  removing _Igeoreg_6
p = 0.9454 >= 0.0001  removing _Iroof2_4
p = 0.9397 >= 0.0001  removing _Iwater2_3
p = 0.8295 >= 0.0001  removing _Igeoreg_3
p = 0.8364 >= 0.0001  removing _Igeoreg_2
p = 0.8528 >= 0.0001  removing morc
p = 0.7703 >= 0.0001  removing _Iqtland_5
p = 0.7612 >= 0.0001  removing _Iwater2_6
p = 0.6986 >= 0.0001  removing _Ifuel3_4
p = 0.6582 >= 0.0001  removing _Iwater2_5
p = 0.6325 >= 0.0001  removing _Iroof2_3
p = 0.4720 >= 0.0001  removing _Iroof2_2
p = 0.2444 >= 0.0001  removing _Ifloor2_4
p = 0.5697 >= 0.0001  removing _Ifloor2_2
p = 0.6573 >= 0.0001  removing _Ifloor2_3
p = 0.2262 >= 0.0001  removing _Iedcsp2_1
p = 0.1688 >= 0.0001  removing _Iedcsp2_4
p = 0.1529 >= 0.0001  removing _Iwater2_4
p = 0.1485 >= 0.0001  removing _Iedcsp2_5
p = 0.1833 >= 0.0001  removing _Iedcsp2_2
p = 0.2874 >= 0.0001  removing _Iedcsp2_3
p = 0.1182 >= 0.0001  removing bicycle
p = 0.1359 >= 0.0001  removing _Ifloor2_5
p = 0.0316 >= 0.0001  removing _Iqtland_2
p = 0.0153 >= 0.0001  removing _Iroof2_5
p = 0.0147 >= 0.0001  removing _Iedchd2_1
```

p = 0.4607 >= 0.0001 removing _Iedchd2_2
p = 0.6762 >= 0.0001 removing _Iedchd2_3
p = 0.5230 >= 0.0001 removing _Iedchd2_4
p = 0.0215 >= 0.0001 removing _Iedchd2_5
p = 0.0027 >= 0.0001 removing toilet2
p = 0.0032 >= 0.0001 removing _Igeoreg_4
p = 0.0346 >= 0.0001 removing _Igeoreg_5
p = 0.0002 >= 0.0001 removing _Iqtland_4
p = 0.0003 >= 0.0001 removing _Iqtland_3
p = 0.0001 >= 0.0001 removing _Ifuel3_6

Probit estimates Number of obs = 3464
LR chi2(10) = 1522.47
Prob > chi2 = 0.0000
Log likelihood = -1621.7678 Pseudo R2 = 0.3194

Allpoor	Coef.	Std. Err.	z	P>|z|	[95% Conf.	Interval]
Carmb	-1.356673	.098397	-13.79	0.000	-1.549527	-1.163818
Radio	-.3148183	.0515732	-6.10	0.000	-.4158998	-.2137367
Bwtv	-.4015334	.0614434	-6.54	0.000	-.5219602	-.2811066
_Iwater2_2	-.647549	.0787468	-8.22	0.000	-.80189	-.4932081
_Ifuel3_5	.2351731	.0582565	4.04	0.000	.1209925	.3493537
_Ifloor2_6	.2937547	.0543398	5.41	0.000	.1872506	.4002587
Colortv	-.9197911	.0672189	-13.68	0.000	-1.051538	-.7880444
Children	.2766634	.0210441	13.15	0.000	.2354176	.3179091
ethnic98	.7469564	.0722465	10.34	0.000	.6053558	.888557
Numfem	.1795615	.0221259	8.12	0.000	.1361955	.2229276
_cons	-.8014307	.0794618	-10.09	0.000	-.957173	-.6456884

. xi: svyprobit allpoor children numfem ethnic i.water2 i.floor2 i.fuel3 radio bwtv colortv carmb if urban92==0
i.water2 _Iwater2_1-6 (naturally coded; _Iwater2_1 omitted)
i.floor2 _Ifloor2_1-6 (naturally coded; _Ifloor2_1 omitted)
i.fuel3 _Ifuel3_3-6 (naturally coded; _Ifuel3_3 omitted)

Survey probit regression

pweight: hhsizewt Number of obs = 4379
Strata: reg10o Number of strata = 7
PSU: commune Number of PSUs = 142
Population size = 59950309
F(20, 116) = 31.83
Prob > F = 0.0000

Allpoor	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
Children	.2722601	.0218202	12.48	0.000	.2291065	.3154137
Numfem	.192731	.0269894	7.14	0.000	.1393544	.2461077
ethnic98	.7352051	.1287394	5.71	0.000	.4805982	.9898121
_Iwater2_2	-.4077328	.3879911	-1.05	0.295	-1.17506	.3595941
_Iwater2_3	.1821916	.3922842	0.46	0.643	-.5936258	.9580089
_Iwater2_4	.0301333	.4061266	0.07	0.941	-.7730601	.8333267
_Iwater2_5	.1523806	.3930176	0.39	0.699	-.6248872	.9296484
_Iwater2_6	.2199634	.4252101	0.52	0.606	-.6209714	1.060898
_Ifloor2_2	.2208648	.1658668	1.33	0.185	-.1071687	.5488983
_Ifloor2_3	.2394117	.1559327	1.54	0.127	-.0689753	.5477986
_Ifloor2_4	.2505421	.1779964	1.41	0.162	-.1014801	.6025642
_Ifloor2_5	.38804	.2394005	1.62	0.107	-.0854204	.8615005

_Ifloor2_6	.5132126	.1626936	3.15	0.002	.1914548	.8349703
_Ifuel3_4	.1208914	.2705578	0.45	0.656	-.4141886	.6559714
_Ifuel3_5	.7872271	.2568654	3.06	0.003	.2792264	1.295228
_Ifuel3_6	.5165502	.2615329	1.98	0.050	-.0006813	1.033782
Radio	-.2941466	.058698	-5.01	0.000	-.4102332	-.17806
Bwtv	-.3936239	.069086	-5.70	0.000	-.5302547	-.2569931
Colortv	-.8476835	.0725484	-11.68	0.000	-.9911619	-.704205
Carmb	-1.356951	.1115015	-12.17	0.000	-1.577467	-1.136436
_cons	-1.754305	.4479311	-3.92	0.000	-2.640175	-.8684351

. gen earthfloor=0
. replace earthfloor=1 if floor2==6
(1829 real changes made)

. gen fuel56=0
. replace fuel56=1 if fuel3==5|fuel3==6
(4423 real changes made)

. xi: svyprobit allpoor children numfem ethnic earthfloor fuel56 radio bwtv colortv carmb if urban92==0

Survey probit regression

pweight:	hhsizewt	Number of obs	=	4381
Strata:	reg10o	Number of strata	=	7
PSU:	commune	Number of PSUs	=	142
		Population size	=	59970186
		F(9, 127)	=	65.50
		Prob > F	=	0.0000

Allpoor	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
children	.2741569	.0214673	12.77	0.000	.2317013	.3166125
numfem	.172756	.0267244	6.46	0.000	.1199034	.2256086
ethnic98	.7400842	.1177045	6.29	0.000	.507301	.9728674
earthfloor	.1961043	.0724632	2.71	0.008	.0527944	.3394142
fuel56	.6347468	.1154854	5.50	0.000	.4063523	.8631414
radio	-.3214317	.0581073	-5.53	0.000	-.43635	-.2065133
bwtv	-.4405607	.0674797	-6.53	0.000	-.5740148	-.3071067
colortv	-.8659367	.0681686	-12.70	0.000	-1.000753	-.7311203
carmb	-1.352716	.1098062	-12.32	0.000	-1.569878	-1.135553
_cons	-1.349279	.124026	-10.88	0.000	-1.594565	-1.103994

. predict cpirur if urban92==0
(option p assumed; Pr(allpoor))
(1618 missing values generated)

. roctab allpoor cpirur [fw=hhsizewt] if urban92==0, graph summary

Obs	ROC Area	Std. Err.	-Asymptotic Normal- [95% Conf.	Interval]
59970186	0.8473	0.0000	0.84723	0.84742

References

Baulch, B. and Hoddinott, J., 2000, *Economic Mobility and Poverty Dynamics in Developing Countries*, London and Portland: Frank Cass

Besley, T. and Kanbur, R., 1993, 'The Principles of Targeting, in M. Lipton and J. Van der Gaag (eds), *Including the Poor*, Washington, D.C.: World Bank Regional and Sectoral Studies

Conway, T., 2001, 'Using government data to target activities to poor communes and monitor poverty reduction: a review of options for the Cao Bang – Bac Kan Rural Development Project', mimeo, Commission of the European Communities, Hanoi

Cornia, G. and Stewart, F., 1995, 'Two Errors of Targeting' in D. Van de Walle and K. Nead (eds), *Public Spending and the Poor*, Baltimore and London: Johns Hopkins University Press

Devereux, S., 2001, 'Food Security Information Systems' in S. Devereux and S. Maxwell (eds), *Food Security in Sub-Saharan Africa*, London: ITDG Publishing

Fofack, H., 2000, 'Applying Household Expenditure Survey Data to Improve Poverty Targeting: the case of Ghana', in D. Bigman and H. Fofack (eds), *Geographical Targeting for Poverty Alleviation: methods and applications*, Washington, D.C.: World Bank Regional and Sectoral Studies

Gelbach, J. and Pritchett, L., 2000, 'Indicator targeting in a political economy: leakier can be better', *Journal of Policy Reform*, Vol 4 No 2: 113–45

Gibson, J., 2001, 'Measuring chronic poverty without a panel', *Journal of Development Economics*, Vol 65: 243–66

Glewwe, P, 1992, 'Targeting assistance to the poor: efficient allocation of transfers when household income is not observed', *Journal of Development Economics*, Vol 38: 297–321

Glewwe, P. and Yansaneh, I, 2000, 'The development of future households surveys in Vietnam', mimeographed report, United Nations Development Program and World Bank, Hanoi

Grosh, M., 1995, 'Towards Quantifying the Trade-Off: administrative costs and incidence in targeted programs in Latin America', in D. Van de Walle and K. Nead (eds), *Public Spending and the Poor,* Baltimore and London: Johns Hopkins University Press

Grosh, M. and Baker, J., 1995, 'Proxy means tests for targeting social programs: simulations and speculation', *Living Standards Measurement Survey Working Paper* 118, Washington, D.C.: World Bank

Haughton, J., 2000, 'Ten puzzles and surprises: economic and social change in Vietnam, 1993–1998', *Comparative Economics Studies,* Vol 42 No 4: 67–92

International Monetary Fund and World Bank, 2001, *Vietnam: Assessment of Interim-Poverty Reduction Strategy Paper*, Washington, D.C: International Monetary Fund and World Bank

MFPED, 2001, 'Poverty indicators in Uganda', *Discussion Paper* 4, Uganda: Ministry of Finance, Planning and Economic Development

Minot, N., and Baulch, B., 2002, 'The spatial distribution of poverty in Vietnam and the potential for targeting', *Policy Research Working Paper* 2829, Washington, D.C.: World Bank

Poverty Working Group, 1999, *Vietnam: attacking poverty*, a joint report of the Government of Vietnam-Donor-NGO Poverty Working Group presented to the Consultative Group Meeting for Vietnam

Rahman, H.Z. and Hossain, M., 1995, *Rethinking Rural Poverty: Bangladesh as a case study*, New Dehli: Sage Publications

Ravallion, M., 1994, *Poverty Comparisons*, Fundamentals of Pure and Applied Economics, Vol 56, Chur: Harwood Academic Publishers

Ravallion, M. and Chao, K., 1989, 'Targeted policies for poverty alleviation under imperfect information: algorithms and applications', *Journal of Policy Modeling*, Vol 11 No 2: 213–24

Sen, A.K., 1985, *Commodities and Capabilities,* Amsterdam: North Holland

Stata Corporation, 2001, 'Receiver operating characteristics (ROC) analysis', *Stata 7 Reference Manual,* Vol 3: 131–51, College Station, Texas: Stata Press

Van de Walle, D., 1995, 'Incidence and Targeting: an overview of implications for research and policy' in D. Van de Walle and K. Nead (eds), *Public Spending and the Poor: theory and evidence,* Baltimore and London: Johns Hopkins University Press

Wodon, Q., 1997, 'Targeting the poor using ROC curves', *World Development,* Vol 25 No 12: 2083–92

World Bank, 1999, *CWIQ (Core Welfare Indicators Survey) Handbook*, Washington, D.C.: African Operational Quality and Knowledge Service, World Bank

Zeller, M., Sharma, M., Henry, C. and Lapenu, C., 2001, 'An operational tool for evaluating the poverty outreach of development policies and projects', *Food Consumption and Nutrition Division Discussion Paper* 111, Washington, D.C.: International Food Policy Research Institute