



United Nations
University

WIDER

World Institute for Development Economics Research

Research Paper No. 2004/48

Commune-Level Estimation of Poverty Measures and its Application in Cambodia

Tomoki Fujii *

July 2004

Abstract

In this study, we combined the Cambodian socioeconomic survey for 1997 and the country's population census of 1998 to produce poverty measures at the commune-level in Cambodia using the small-area estimation technique developed by Elbers, Lanjouw and Lanjouw. While there are a number of communes for which the standard errors associated with the estimates are too high, we found that they are on average small enough to be useful. We illustrate the application of poverty maps with reference to education programmes.

Keywords: poverty mapping, Cambodia, targeting, education

JEL classification: O20, I28, I32

Copyright © UNU-WIDER 2004

* University of California, Berkeley.

This study has been prepared within the UNU-WIDER project on Spatial Disparities in Human Development, directed by Ravi Kanbur and Tony Venables, with Guanghai Wan.

UNU-WIDER gratefully acknowledges the financial contributions to its research programme by the governments of Denmark (Royal Ministry of Foreign Affairs), Finland (Ministry for Foreign Affairs), Norway (Royal Ministry of Foreign Affairs), Sweden (Swedish International Development Cooperation Agency–Sida) and the United Kingdom (Department for International Development).

ISSN 1810-2611

ISBN 92-9190-637-9 (internet version)

Acknowledgements

This paper is a revised version of the report titled ‘Estimation of Poverty Rates at Commune-Level in Cambodia: Using the Small-Area Estimation Technique to Obtain Reliable Estimates’, prepared for the World Food Programme. Without implicating, the author thanks Rebecca Hansen, Jean Olson Lanjouw, Peter Lanjouw, Livia Montana, Berk Özler, Mahadevan Ramachandran, H. E. Kim Saysamalen, Boreak Sik, two anonymous referees and the participants of the UNU-WIDER Conference on Spatial Inequality in Asia held in Tokyo in March, 2003. The views expressed in this paper do not necessarily represent those of the World Food Programme. All the remaining mistakes are mine.

The World Institute for Development Economics Research (WIDER) was established by the United Nations University (UNU) as its first research and training centre and started work in Helsinki, Finland in 1985. The Institute undertakes applied research and policy analysis on structural changes affecting the developing and transitional economies, provides a forum for the advocacy of policies leading to robust, equitable and environmentally sustainable growth, and promotes capacity strengthening and training in the field of economic and social policy making. Work is carried out by staff researchers and visiting scholars in Helsinki and through networks of collaborating scholars and institutions around the world.

www.wider.unu.edu

publications@wider.unu.edu

UNU World Institute for Development Economics Research (UNU-WIDER)
Katajanokanlaituri 6 B, 00160 Helsinki, Finland

Camera-ready typescript prepared by Liisa Roponen at UNU-WIDER
Printed at UNU-WIDER, Helsinki

The views expressed in this publication are those of the author(s). Publication does not imply endorsement by the Institute or the United Nations University, nor by the programme/project sponsors, of any of the views expressed.

1 Introduction

Cambodia is still suffering from the legacy of civil conflict after more than a decade. With over a third of the population living under the poverty line, poverty remains one of the most serious problems in Cambodia. A number of governmental bodies, local and international NGOs, and international organizations operating in Cambodia have made eradicating poverty a priority and have established many social programmes to this end.

In designing such programmes, efficient allocation of resources is essential for making poverty alleviation more cost-effective. Targeting is often helpful for this purpose because one can avoid wasting resources on the non-poor, which would occur in the absence of targeting. If targeting were costless and consumption poverty were at issue, it would be desirable to formulate such a targeting policy that the gap between current consumption and the poverty line were just fulfilled.

In reality, such a social intervention is unlikely to be possible for at least two reasons. First, targeting is not costless as there are administrative, political and other costs. Van de Walle (1998) argues that hidden costs associated with targeting may undermine its benefits and these need to be evaluated. Second, there are two types of errors in targeting. One is the error of exclusion, in which intended beneficiaries cannot benefit from the intervention. The other is the error of inclusion, in which an intervention reaches individuals who were not intended to be beneficiaries (Hoddinott 1999). These errors are commonly known as Type I and Type II errors, respectively and, usually, these two error types cannot be eliminated altogether.

The question policymakers face is whether and what type of targeting policy would make anti-poverty programmes more efficient. One of the obvious ways of targeting is to move resources according to geographic information. Such targeting is called geographic targeting. Geographic targeting is relatively easy to monitor and administer, and local institutions and NGOs can greatly assist in implementing the programmes. Moreover, geographic targeting has relatively little influence on a household's behaviour because it is difficult and costly for a household to change its place of residence (Bigman and Fofack 2000a). Furthermore, geographic targeting can be combined with other targeting mechanisms or criteria. For example, a food-for-work programme for local road construction is a combination of geographic targeting and self-targeting. Such a programme may be best targeted to poor areas with insufficient road access.

As a powerful tool to identify the location of the poor and to enable a more efficient allocation of resources through geographic targeting, poverty mapping has increasingly drawn the attention of researchers and practitioners in the development community. Poverty maps may be combined with other maps to derive even more valuable information by overlaying maps using geographic information systems (GIS). For example, the targeting policy of basic medical care programme may be better formulated by choosing poor areas where indicators show low health status. Irrigation systems may be most beneficial in poor areas that are drought-prone. Such targeting would be difficult or impossible without poverty maps.

This paper has two purposes. The first purpose is to present in the form of maps the commune-level estimates of poverty measures using the small-area estimation technique recently developed by Elbers *et al.* (2000, 2002, 2003), which allows for explicit

treatment of the standard errors. The maps are helpful in identifying target areas and formulating efficient and effective programmes and policies to reduce poverty. Furthermore, they provide a basis for promoting coordination within and between donors and implementing organizations in addressing poverty.

The second purpose of this paper is to illustrate its application. Though poverty maps *per se* are useful to formulate targeting policies, it is not clear whether a certain type of social-sector intervention programme is suitable for certain areas. Using the school feeding-programme as an example, we shall discuss how best other maps may be combined with the poverty maps to identify the target areas for social-sector intervention programmes.

This paper is organized as follows. Section 2 reviews the relevant literature. Section 3 explains the theory of poverty mapping. Section 4 describes the data used in this study. Some known problems with the dataset are also discussed. It should be noted that the extensive use of geographic data characterizes this study. The implementation of poverty mapping is provided in section 5, followed by the results in section 6. Section 7 illustrates the application of poverty mapping to school feeding-programmes and section 8 concludes.

2 Review of literature

2.1 Motivation of poverty mapping

Socioeconomic surveys have been widely used to analyse poverty in developing countries. Cambodia is no exception. Three successive Cambodian socioeconomic surveys (CSES), carried out in 1993/4, 1997 and 1999, respectively, have been used to generate poverty estimates. They have provided valuable information for policymakers (Prescott and Pradhan 1997; MoP 1998, 2001). However, the sampling designs of these surveys impose severe limitations on the geographic level of disaggregation at which the poverty estimates are reliable. None of the three surveys can provide a reliable estimate of poverty rates even at the provincial level.¹

However, as is often the case, what policymakers really need is information that is geographically disaggregated. They may want poverty estimates at the district or even commune level. If policymakers want to deliver food to poor people, knowing poverty rates at the provincial level may not be very useful, as too many non-poor may benefit due to the error of inclusion. Baker and Grosh (1994) examined the efficiency gains from targeting in Mexico, and showed that only small improvements over uniform transfers of money can be made if the programme is targeted at the state level, but the improvement is significant if targeted at the district or neighbourhood level.

Whether detailed targeting is possible depends on the information available to policymakers. Poverty maps allow reducing informational constraints, which are one of the central issues concerning the formulation of targeting policies (Ravallion and Chao 1989; Kanbur 1987). They help policymakers to find pockets of poverty, or poor areas

¹ In Cambodia, there are four levels of administrative divisions, called province, district, commune and village in the descending order of aggregation.

surrounded by non-poor areas, which cannot be identified from the socioeconomic surveys alone. Hence, poverty maps are useful for formulating geographic targeting policies to move assistance to the neediest people in a more efficient and transparent manner.

2.2 Household unit small-area estimation

Poverty maps can be created by a number of methodologies including the small-area estimation, multivariate weighted basic needs index, combination of qualitative information and secondary data, and extrapolation of participatory approaches. Davis (2002) overviews these various poverty mapping methods and their applications, and discusses their merits and limitations. Small-area estimation is a statistical technique that combines survey and census data to derive statistics for geographically small areas such as communes and districts.

There are two variants called the household unit-level method and the community-level data method, depending on the level at which the census records are available. The basic idea for these methods is that the welfare measure at the household level or community level is regressed on a set of variables that are common between the census and the socioeconomic survey. Then the welfare measure is imputed in each record in the census. The advantage of running regression at the household level is that the standard errors associated with poverty estimates can be evaluated through regression, while it is often easier to access the community-level census data, and computational burden is substantially lower. Vietnam has a poverty map based on the community-level data method (Minot 2000). Other examples include Bigman *et al.* (2000) for Burkina Faso, and Bigman and Fofack (2000b) for India.

The household unit-level method was first applied to Ecuador (Hentschel *et al.* 2000). Its statistical properties were rigorously studied and various estimation strategies were discussed by Elbers, Lanjouw and Lanjouw (2003, hereafter ELL). The ELL approach has been applied to a number of countries. Alderman *et al.* (2002) studied the case in South Africa and found that the income from census data provides only a weak proxy for the average income or poverty rates at either the provincial level or at lower levels of aggregation. Demombynes *et al.* (2002) compared the experience of poverty mapping in Ecuador, Madagascar and South Africa. Cambodia is the first country in Asia to which the ELL technique is applied.

3 Methodology of poverty mapping

To measure the welfare of people, we use household consumption expenditure, defined as goods and services bought on the market, received in kind, or produced by the household. While the consumption measure captures only certain aspects of poverty, it is a commonly used and widely studied welfare measure, and has a relatively straightforward interpretation. All of the consumption items in the CSES questionnaire, including food items and non-food items, are aggregated to arrive at the consumption aggregate for the household. When the *per capita* consumption is derived, the household consumption is simply divided by the number of individuals in the household. We report the Foster-Greer-Thorbecke (FGT) poverty measures with parameter 0, 1 and 2, which are commonly called poverty rate, poverty gap and poverty

severity. The poverty lines used in this study to derive 5 FGT poverty measures are given in section 5.

The concepts behind poverty mapping are straightforward. First, the survey data are used to estimate a consumption model. This model describes the relationship between consumption and right-hand side variables. Right-hand side variables are restricted to those variables that can also be found in the census or in a tertiary dataset that can be linked to both the census and the survey. A geographic information system (GIS) dataset is used as tertiary data. The census data are then fed into the model with the parameter estimates to derive the small-area statistics of interest. It should be noted that we assume the models estimated from the survey data to be applicable to census records. The theoretical underpinnings of this methodology are given in detail in a series of papers by Elbers *et al.* (2000, 2002, 2003). One of the important contributions of their study is that they rigorously examined the properties of the small-area statistics. In what follows, we shall briefly present the feature of the theory. *Per capita* household consumption, y_h , for household h is related to a k -vector of observable characteristics, \mathbf{x}_h , through the following model.

$$\ln y_h = \mathbf{x}_h^T \boldsymbol{\beta} + u_h \quad (1)$$

where $\boldsymbol{\beta}$ is a k -vector of parameters and u_h is a disturbance term. u_h satisfied $E[u_h | \mathbf{x}_h] = 0$. As described in section 5, the disturbance term is decomposed into the location, or cluster-specific, effect and the household-specific effect to allow for spatial autocorrelation and heteroscedasticity among households in application. The parameter $\boldsymbol{\beta}$ is estimated through regression using the household survey data. The empirical distribution of the residual terms is also obtained from the regression. This regression will be referred to as the first-stage regression.

For the purpose of poverty maps, the consumption of each household is not of interest but rather the various welfare measures at a certain level of aggregation. In this paper, commune-level aggregation was chosen because such a level of aggregation is *useful* and the estimate at that level is *acceptable*. Welfare estimates at a more aggregated level such as the district or provincial level are more accurate, and depending on the purpose, accuracy may be more important than the level of disaggregation.

The problem we face in this exercise is that we only know \mathbf{x}_h^T , but would like to estimate a commune-level welfare measure W derived from y_h for households in the commune. A natural way to compute W is to impute y_h using the parameters estimated in the first-stage regression and use the imputed values. Such an estimate is subject to statistical errors. Elbers *et al.* (2002) showed that the error can be decomposed into the idiosyncratic error, model error and computation error, and studied their properties.

The magnitude of idiosyncratic error is approximately in inversely proportional to the size of the population in the target area. Hence, there is a practical limit to the degree of disaggregation possible. This is precisely the reason why we did not produce the estimates at the village level. The model error in general depends solely on the standard errors of the first-stage coefficients and the sensitivity of the indicators to deviations in household consumption. The computation error depends upon the computational method used. Using simulation methods with sufficient computational resources and time, this error can be made arbitrarily small.

4 Datasets

4.1 The CSES datasets

To produce the poverty maps, we used four distinct datasets—two socioeconomic survey datasets, a census dataset and a GIS dataset. The consumption model is built upon the two socioeconomic surveys, namely, the CSES 1997 and the CSES 1999. For reasons discussed later, the CSES 1999 was used only for auxiliary purposes. The CSES 1997 was conducted by the National Institute of Statistics of the Ministry of Planning (MoP) and funded by UNDP and the Swedish International Development Agency, and executed by the World Bank. The questionnaires for CSES 1997 included three substantive components: a village questionnaire, a core questionnaire for households, and a social sector household module. Only the core questionnaire was used in this study.

The sample design for the CSES 1997 treated villages as the primary sampling units and households as secondary sampling units. A sampling frame developed for the Socioeconomic Survey of Cambodia 1996 was updated with newly available information to use as the sampling frame for the CSES 1997 survey. In the CSES 1997, there were three sampling strata: Phnom Penh, ‘other urban’, and ‘rural’. The total sample size of the CSES 1997 was 6,010 households in 474 villages. In the Phnom Penh stratum, a sample was taken from 120 sample villages with 10 households from each village, while in the ‘other urban’ stratum, 10 households from each of 100 villages were sampled. In the ‘rural’ stratum, 15 households were sampled in each of 254 villages. For each of the three sampling strata, a consumption model for small-area estimation was constructed.²

The CSES 1999 is similar to the CSES 1997 in design, and was carried out in two rounds between January and March 1999 and between June and August 1999 to capture seasonal variations in consumption. While the CSES 1999 also had three components, we only used the core questionnaire. The CSES 1999 has ten sampling strata defined from the urban and rural sectors within each of five zones (Phnom Penh, Plain, Tonle Sap, Coastal and Plateau).³ Though the CSES 1999 is more attractive for our purposes in terms of the sampling design than the CSES 1997, the inconsistencies in measured consumption between the two rounds of the survey indicate the potential presence of widespread and systematic measurement error (MoP 2001). Since the reliability of the consumption measure is critical in this study, we decided to use the CSES 1997 for the most part of our analyses.

Due to security concerns and access limitation, some parts of Cambodia were not covered in the sampling frame of the CSES 1997 and the CSES 1999. In terms of the number of households, 11.6 per cent of the rural areas and 2.6 per cent of the urban areas were not covered in CSES 1997. The corresponding figures for the CSES 1999 were substantially smaller; 3.8 per cent for the rural areas and 0.3 per cent for the urban areas. Hence, we decided to take advantage of the better geographical coverage of the sampling frame for CSES 1999 to see if the consumption model holds for those

² For further details of CSES 1997, see NIS (1998).

³ For further details of CSES 1999, see NIS (2000a).

households outside the CSES 1997 sampling frame, but inside the CSES 1999 sampling frame.

To do so, we took the following steps. First, using data from NIS (1997), we identified the villages in the CSES 1999 that were excluded from the CSES 1997 sampling frame. Then we estimated the parameters of the consumption model using the CSES 1999 data. Second, we assigned each record in the CSES 1999 dataset the corresponding stratum code of the CSES 1997. Ideally, two regressions with the same set of regressors should be run separately for the areas inside and outside the sampling frame of CSES 1997 to test the hypothesis that the estimated parameters for those two areas are the same. However, the sample sizes for the excluded areas were too small to allow one to generate meaningful results. Instead, we took an alternative approach. We first ran a regression without the excluded areas and the coefficient β_0 was estimated. Then we ran another regression with the excluded areas and estimated the coefficient β_1 . We tested the hypothesis $\beta_0 = \beta_1$. The rejection of this hypothesis would suggest that, if the CSES 1997 had included the excluded areas, the consumption model would have been different.

It should be noted that in the procedure described above, the CSES 1999 dataset does not affect the estimated parameters used in the simulation. This is because there is concern about the quality of the CSES 1999 dataset. However, as was observed in MoP (2001), there were some common patterns between the two rounds. This seems to suggest that the overall pattern of consumption was not altered to the extent to render the test described above invalid. The tests were carried out with CSES 1999 round 1, CSES 1999 round 2 and both rounds pooled.

4.2 The census dataset

The Cambodian national population census was conducted over a period of ten days in March 1998. It was the first population census to be conducted in Cambodia since 1962 and was done on a *de facto* basis. The census covered all persons staying in Cambodia, including foreigners, at the reference time, which was the midnight of 3rd March 1998. Foreign diplomatic corps and their families were, however, excluded. The census questionnaire consisted of two forms, form A (the house list) and form B (the household questionnaire). The construction material of walls, roof and floors of each house was observed by the enumerator and recorded in form A together with other information. Form B had four parts; part 1 collected information on the usual household members present and absent on census night as well as visitors present on the census night; part 2 gathered specific information on each usual household member and visitors present on the reference night, including full name, relationship to the head of household, sex, age, marital status, mother tongue, religion, birth place, migration, literacy, education and employment. Part 3 contained questions on fertility of females ages 15 and over, and part 4 contained housing characteristics, conditions and other facilities.

The geographical frame for the census followed the defined structure of province, district, commune, and village in descending order of aggregation. There are 24 provinces in Cambodia, including the municipality of Phnom Penh, and the towns of Kep, Sihanoukville, and Krong Pailin. Due to military operations, fifteen communes with an estimated population of 45,000 were not covered in the census. Since it is not

possible to estimate poverty measures for these areas, we analysed only the 1,594 communes included in the census.⁴

Prior to the computer simulation, we applied two treatments to the census dataset. First, we excluded special settlements from the dataset. Special settlements are groups of people who were found together on the census night. These people are transitional and may not necessarily live in the commune. Hence, they were not included for the calculation of the poverty estimates.

Second, there is a practical inconsistency between the definitions of household used in the census and survey datasets. Even though the census dataset distinguishes between usual members of household and visitors in form A, part 2 of form B includes both if they were present on the reference night, and makes no distinction between them. This means that the data user has to take the usual members of the household as well as visitors present on the census night as the household. The survey, however, asks questions about the usual members of the household, including those absent at the time of survey. Moreover, there were households that did not appear to be regular households. For example, there were households with more than 100 people.

Hence a practical decision was made to take care of this issue. Only those households for which the household size is less than 16 and the number of visitors is less than 10 were used for the analysis.⁵ The original dataset contained 2,162,086 regular households and, as a result of this treatment, was reduced to 2,150,235 households. Admittedly, the decision may have been somewhat arbitrary. However, it seemed more reasonable to make such a distinction than to ignore the issue. More importantly, the exclusion does not affect the main findings of this study significantly because much less than 1 per cent of the dataset was dropped.

4.3 The GIS dataset

A set of geographically derived indicators was also used in this analysis. These indicators included distance calculations, land use and land cover information, climate indicators, vegetation, agricultural production, and flooding. A number of datasets from various sources were compiled into a GIS and the geographic indicators were generated for all villages and communes in Cambodia. Very coarse resolution data were summarized at the commune level, while high resolution data were attributed to individual villages. Distances from villages to roads, other towns, health facilities, and major rivers were calculated. Indicators, based on satellite data with varying temporal resolutions, included land use within the commune (agricultural, urban, forested, etc.), a vegetation greenness indicator to proxy agricultural productivity, and the degree to which the area benefited from nighttime lighting as a proxy of urbanization. Relatively stable indicators, including soil quality, elevation, and various 30-year average climatological variables, were also generated from other composite datasets.

⁴ Further information on the implementation of the census can be found in NIS (2000b).

⁵ The maximum household size is 14 in CSES 1997 data. We allowed for a margin of one. CSES 1997 does not contain information on visitors. We excluded those households with clearly too many visitors.

The sources as well as the spatial and temporal dimensions of the datasets vary. Some datasets were assumed to not have changed greatly over time. Others, where multi-temporal data were available, included both yearly and monthly indicators as well as change and long-term average indicators. Road, river, village location, and administrative boundary data were obtained under a UNTAC project and updated in 1996 by the Department of Geography under a UNDP-sponsored CAREERE project. Health facility latitude and longitude locations were provided by the World Health Organization. Land use and land cover data were obtained from Landsat Thematic Mapper satellite for 1993 and 1997 at 50 meter resolution. Agricultural production data at the commune-level were taken from the commune-level crop assessment database prepared by World Food Programme (WFP). NASA's Advanced Very High Resolution Radiometer (AVHRR) satellite data at 7km resolution were used to generate the normalized differential vegetation index (NDVI). A 19-year monthly series of AVHRR-derived NDVI data, covering 1981-2000 and compiled by Clark Labs, were used to generate the NDVI values. NDVI indicators included monthly values, 19-year average and standard deviation, and coefficient of variation.

A global digital elevation model at 1km resolution, GTOPO30, was used for elevation values. GTOPO30 was developed under the coordination of the US Geological Survey, in collaboration with NASA, UNEP/GRID, USAID and others. City lights satellite data at 1km resolution were collected during 1994-95 by the Defense Meteorological Satellite Programme and obtained from the National Geophysical Data Center. The soil-quality data are based on a reclassification of the FAO/UNESCO Soils Map of the World, which contains 106 soil type classes. The USDA Natural Resources Conservation Service and the University of Puerto Rico overlaid the FAO/UNESCO map with a global climate dataset, and using the combined climate and soils data, reclassified the FAO/UNESCO soils map according to suitability for food production. The University of East Anglia Climatic Research Unit's Global Climate dataset was obtained from the Intergovernmental Panel on Climate Change. These 30-year monthly averages, interpolated into five degree grids, are based on daily weather-station data collected from 1961 to 1990.

5 Implementation of the poverty mapping

5.1 The choice of consumption aggregate and poverty line

Using the CSES 1997 to define the consumption aggregate was not as straightforward as it initially seemed. There were two possible alternatives. One alternative was to use the adjusted consumption aggregate derived by Knowles (1998) while the other alternative was the unadjusted consumption aggregate defined in MoP (2001). When MoP (1998) was published, the dataset contained errors which necessitated the use of the adjusted consumption aggregate. The mistakes were subsequently corrected and hence the adjustments made by Knowles are unnecessary for this paper. Therefore, this paper follows the definition of unadjusted consumption given in MoP (2001).

To ensure comparability with the publicized benchmark national poverty rate of 36.1 per cent, we redefined the poverty line so that the same poverty rates could be reproduced using the unadjusted consumption aggregate for each of the three strata. As a result, the poverty lines, in terms of *per capita* per day consumption, employed in this

analysis are 1,629 riels for Phnom Penh, 1,214 riels for other urban and 1,036 riels for rural.⁶

By construction, the poverty rate for each stratum in this paper is the same as given in MoP (1998). However, there is no guarantee that the poverty gap and the poverty severity are the same. Hence checking these indices provides an indication of how important the choice between the adjusted and unadjusted measures is. In this paper, the poverty gap was estimated at 8.9 per cent and the poverty severity at 3.2 per cent. The corresponding figures in MoP (1998) were 8.7 per cent and 3.1 per cent, respectively. Although these numbers are not exactly the same, the differences are small enough to be considered random errors. This seems to suggest that the analysis presented in this paper will be robust with respect to the choice of consumption aggregate.

5.2 The consumption model

This section provides details of the implementation of the ELL approach in Cambodia. We followed the implementation described in Elbers *et al.* (2002) whenever possible. As was discussed above, the CSES 1997 has three strata and is intended to be representative at that level. Thus, we have constructed three consumption models, one for each stratum. Hereafter, subscripts v and h are used to denote a village and a household, respectively.

The first step in creating a poverty map is developing an accurate empirical model of household consumption.⁷ It should be noted that our interest here is not in the description of causal relationship, but in the prediction of household consumption. We estimate the following consumption model:

$$\ln y_{vh} = E[y_{vh} \mid \mathbf{x}_{vh}^T] + u_{vh} = \mathbf{x}_{vh}^T \boldsymbol{\beta} + \eta_v + \varepsilon_{vh} \quad (2)$$

where y_{vh} is the *per capita* consumption and u_{vh} is the disturbance term, which is the sum of the location-specific component η_v and the idiosyncratic component ε_{vh} . These two components, η_v and ε_{vh} , are assumed to be independent of each other and uncorrelated with observable household characteristics \mathbf{x}_{vh} . This specification allows for an intra-cluster (i.e., intra-village) correlation in the disturbances and heteroscedasticity in ε_{vh} . Explicit treatment of the location effects is important as some of the effects of location may remain unexplained even with a rich set of regressors, including a number of geographic variables. The household characteristics \mathbf{x}_{vh} in this model are not limited to variables that are specific to the household. They can also include the characteristics of the village in which the household is located. For example, \mathbf{x}_{vh} can include the village-level means of the census data and the GIS data, which capture a part of the location effects. Crossterms between a household-level variable and a GIS variable were also included.⁸ For notational convenience, the variance of a random variable will be

⁶ The average official exchange rate in 1998 was US\$=3807.8 riel.

⁷ The detailed regression results in this section are omitted to save space. They are given in WFP (2002), and also available from the author upon request.

⁸ While the fact that the cross-terms increased the predictive power of the model justifies their inclusion in the model, it is plausible that the interaction term indeed affects the welfare measure. For example,

hereafter denoted as $\sigma^2 \equiv \text{Var}[\cdot]$. When \cdot has a subscript s , it is expressed using a comma as $\sigma_{\cdot,s}^2 \equiv \text{Var} \left[\begin{matrix} \cdot \\ s \end{matrix} \right]$.

Elbers *et al.* (2002) point out that, for any given disturbance variance $\sigma_{u,vh}^2 = \sigma_{\eta,v}^2 + \sigma_{\varepsilon,vh}^2$, the greater the fraction due to the common component, the less one enjoys the benefits of aggregating over more households within a country. To assess the performance of the consumption model, a number of diagnostic statistics are checked, and important ones are reported in Table 1.⁹ In all the three models, R^2 statistics are reasonably high, with the urban stratum performing the best. Since unexplained location effects reduce the precision of poverty estimates, we tried to explain the variation in consumption due to location as far as possible with the choice and construction of \mathbf{x}_{vh} . Location means of household-level variables derived from the census data are particularly useful for this purpose. The importance of location effect as measured by $\hat{\sigma}_{\eta}^2 / \hat{\sigma}_u^2$ is quite small for the Phnom Penh and other urban strata, but is relatively high in the rural stratum.¹⁰

In constructing the models, we also need to be careful not to overfit the data. While the model error can grow as the number of indicators used in the model increases, overfitting is still a concern. To address this issue, we carried out random dropping and confirmed the robustness of the coefficients, as was done in Elbers *et al.* (2002). Random dropping was tried not only at the individual level but also at the village level, because, when too many village-level indicators are included in the model, we are just fitting the village-level means.¹¹

Table 1
Diagnostic statistics

Description	Phnom Penh	Other urban	Rural
No. of observations	1200	1000	3810
No. of clusters	120	100	254
# HH-level RHS	22	10	17
# census mean vars	9	13	35
First stage # GIS vars	0	5	11
Regression # cross terms b/w HH-level and GIS	7	12	12
R^2	0.481	0.700	0.538
$\hat{\sigma}_{\eta}^2 / \hat{\sigma}_u^2$	0.047	0.008	0.139
Residual # RHS Var	29	23	34
Regression R2	0.121	0.066	0.045

the potential of good soil quality may be more efficiently exploited by those with better education or experience, which results in positive and significant coefficient on the relevant cross-term.

⁹ WFP (2002) provides a complete list of diagnostic statistics we used.

¹⁰ See Elbers *et al.* (2002) for the derivation of $\hat{\sigma}_{\eta}^2 / \hat{\sigma}_u^2$.

¹¹ We used stepwise regression to reduce the number of candidate variables included in the model. We determined the final models taking into consideration the diagnostic statistics and the concern for overfitting.

Once the consumption models are specified, the next step is to estimate each component of the disturbance term. First, the residual term \hat{u} was derived from the OLS regression. The common component η_v was estimated non-parametrically at the average of \hat{u} in the cluster as follows:

$$\hat{u} = \hat{u}_v + (\hat{u}_{vh} - \hat{u}_v) = \hat{\eta}_v + e_{vh} \quad (3)$$

where \hat{u}_v is the average of \hat{u} over the households in the same village. To model heteroscedasticity in the idiosyncratic part of the residual, a restricted number of household characteristics, z_{vh} , that best explain variation in e_{vh} out of potential explanatory variables, their squares and interactions were chosen. The following logistic model of the variance of ε_{zh} conditional on z_{vh} , bounding the prediction between zero and a maximum, $A \equiv 1.05 \times \max_{v,h} \{e_{vh}^2\}$, was estimated:

$$\ln \left[\frac{e_{vh}^2}{A - e_{vh}^2} \right] = z_{vh}^T \alpha + r_{vh} \quad (4)$$

Letting $B \equiv \exp \{z_{vh}^T \hat{\alpha}\}$ and using the delta method, the model implies a household-specific variance estimator for ε_{vh} of:

$$\hat{\sigma}_{\varepsilon,vh}^2 = \left[\frac{AB_{vh}}{1 + B_{vh}} \right] + \frac{1}{2} \hat{\sigma}_r^2 \left[\frac{AB_{vh}(1 - B_{vh})}{(1 + B_{vh})^3} \right] \quad (5)$$

Once $\hat{\sigma}_{\varepsilon,vh}^2$ is computed, the household residuals are standardized as follows:

$$e_{vh}^* = \frac{e_{vh}}{\hat{\sigma}_{\varepsilon,vh}} - \left[\frac{1}{H} \sum_{vh} \frac{e_{vh}}{\hat{\sigma}_{\varepsilon,vh}} \right] \quad (6)$$

where H is the number of households in the survey. Before proceeding to conduct the simulation, the estimated variance-covariance matrix, $\hat{\Sigma}$, was weighted by the population expansion factor to obtain GLS estimates of the first-stage parameters, $\hat{\beta}_{GLS}$, and their variance $Var(\hat{\beta}_{GLS})$.

5.3 Simulations

From the consumption model, α , β_{GLS} , and their associated variance-covariance matrices as well as the empirical distribution of e_{vh}^* and $\hat{\eta}_v$ are obtained. Assuming multivariate normal distribution, $\tilde{\alpha}^R$ and $\tilde{\beta}^R$ are drawn for each R -th simulation. Once $\tilde{\alpha}^R$ is drawn, the household-specific variance of the household component of disturbance, $(\tilde{\sigma}_{\varepsilon,vh}^2)^R$, is estimated for each census household. Then, the error terms are drawn in two stages to take clustering into account. The location-specific error $\tilde{\eta}_v^R$ is

drawn from the empirical distribution of $\hat{\eta}_v$. Then the household component $\tilde{\varepsilon}_{vh}^R$ is obtained with a draw from the empirical distribution of e_{vh}^* in the corresponding cluster (i.e. village) and $(\tilde{\sigma}_{\varepsilon,vh}^2)^R$. The simulated value of consumption \hat{y}_{vh}^R for household in village is, therefore,

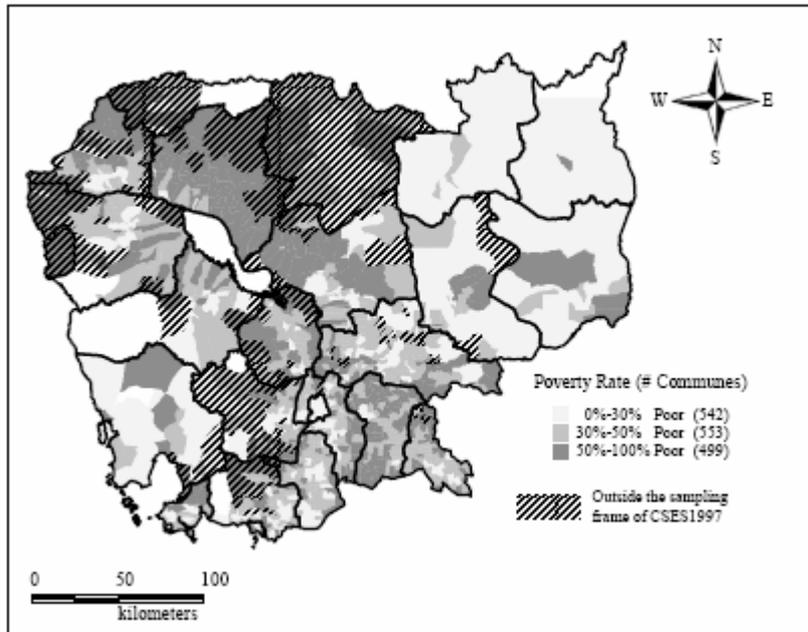
$$\hat{y}_{vh}^R = \exp(\mathbf{x}_{vh}^T \tilde{\beta}^R + \tilde{\eta}_v^R + \tilde{\varepsilon}_{vh}^R) \quad (7)$$

The full set of simulated \hat{y}_{vh}^R is used to compute the R -th estimate of poverty measures for each commune except for some outliers. For example, the R -th estimate of poverty incidence for commune c , \hat{I}_c^R , is computed as follows:

$$\hat{I}_c^R = \frac{1}{n_c} \sum_{v \in V_c} \sum_{h \in H_v} \text{Ind}(\hat{y}_{vh}^R < z) \cdot n_{vh} \quad (8)$$

where V_c denotes the set of villages in commune c , H_v the set of households in village v , n_{vh} the size of household h in village v , z the poverty line, n_c the population of commune c , and $\text{Ind}(\cdot)$ is an indicator function. For this paper, the simulation was repeated 100 times. The mean and standard deviation of the estimates of poverty measures from each simulation were computed to arrive at the commune-level estimates of poverty measures and their associated standard errors. In a similar manner, poverty measures at more aggregated levels, such as district, province and stratum, were estimated.

Figure 1
Commune-level poverty rate of Cambodia



Source: Author

6 Results

6.1 Creating poverty maps

Once the commune-level estimates of poverty measures are computed, it is straightforward to create poverty maps. The polygon data for communes are combined with poverty estimates by the GIS. The map presented in Figure 1 is the map for poverty rate. The darker areas are poorer and the hatched areas are outside the sampling frame of CSES 1997.

6.2 Are the villages excluded from the CSES 1997 sampling frame different?

As noted before, the sampling coverage of the CSES 1997 is smaller than that of the CSES 1999. Hence, we used the CSES 1999 data to check if the consumption model applies to those areas which were excluded from the CSES 1997 sampling frame but included in the CSES 1999 sampling frame. Table 2 provides a summary of the sampling frame of 1997. Unfortunately, the number of samples from the excluded areas was too small to meaningfully compare the equality of the coefficients for the included and excluded areas. Hence, as discussed before, the hypothesis that the estimated regression coefficients with and without excluded areas are the same was tested.

If the relationship between the right-hand side variables and consumption is kept intact in each round, in principle, the same conclusion should be derived. However, the results obtained in this study are mixed. For the ‘other urban’ stratum, the hypothesis was rejected at the significance level of 1 per cent when round-1 data or pooled data (i.e. round 1 and round 2) were used. However, the hypothesis could not be rejected even at the significance level of 10 per cent when round-2 data were used. For the ‘rural’ stratum, the pooled sample could not reject the hypothesis but with the round-1 and round-2 data, both sets rejected it when used separately. Hence the results are inconclusive.

How one should deal with the excluded areas depends on the objective of the policymaker. If one wants to reduce the Type I error and the budget is severely limited, these areas may be best avoided. On the other hand, the policymaker is concerned about lives in the heavily mined areas, which roughly correspond to the excluded areas, they could use the estimates in the initial stages of targeting area selection. Irrespective of

Table 2
Summary of the number of villages in CSES 1999
included and excluded from the sampling frame of CSES, 1997

Stratum	Round 1		Round 2		Total	
	Excluded	Included	Excluded	Included	Excluded	Included
Phnom Penh	0	600	0	600	0	1200
Other urban	50	810	10	850	60	1660
Rural	180	1360	110	1430	290	2790
Cambodia	230	2770	120	2880	450	5650

how the policymaker decides to use the excluded-area estimates, they should bear in mind that there is a good reason for caution in using these estimates.¹²

6.3 How accurate are the estimates?

To have an idea about the accuracy of the estimates, we compared our results with the preliminary poverty map (WFP 2001) constructed by combining survey and census. Because WFP (2001) uses a different combination of datasets from that used in this study and because it does not record standard errors, it is not possible to compare the commune-level estimates directly. It makes more sense to conduct a comparison at the stratum level, because it is possible to use the poverty measures in poverty profiles as benchmarks. Table 3 compares the stratum-level poverty rates obtained in WFP (2001) and in this study for survey data estimates (CSES only) and for estimates obtained by combining the survey and census (CSES plus census). The latter estimates are created at the stratum-level and consistent with the poverty maps. For the sake of comparison, WFP (2001) model 2 was used as it also has a consumption model for each of the three CSES 1997 strata.

Two observations can be made. First, the patterns of poverty rates in this paper and in WFP (2001) are quite similar regardless of the datasets used. The poverty rate in Phnom Penh is around 10-12 per cent, other urban is 25-30 per cent and rural is 40-50 per cent. Second, when the differences between CSES only and CSES plus census for this paper and for WFP (2001) are compared, there are much smaller discrepancies for this paper. The difference between CSES only and CSES plus census for this paper is small enough to be attributed to the random error.

The level of accuracy of the commune-level estimates varies from commune to commune. For example, the standard errors associated with commune-level estimates range 20 between 0.1 per cent and 22.6 per cent. Table 4 provides summary statistics on the standard errors. The first column (mean standard error) is the simple average of the standard errors. Urban areas have lower standard errors. The median of standard errors

Table 3
Stratum level comparison of estimates of poverty rates in percentage

	Stratum	CSES only		CSES plus census	
This study	Phnom Penh	11.1	(1.8)	11.9	(1.2)
	Other urban	29.9	(3.3)	30.0	(1.0)
	Rural	40.1	(2.0)	43.0	(1.1)
WFP (2001) Model 2	Phnom Penh	9.7		12.5	
	Other urban	25.2		26.7	
	Rural	40.1		49.4	

Note: Standard errors in brackets are calculated by the author. The standard errors for CSES Only take into account clustering and expansion factors.

¹² For example, Krong Pailin was the poorest province in the map. It is, however, generally considered to be a non-poor area. This may be because Krong Pailin was out of the sampling frame of the CSES 1997

Table 4
Summary statistics of the standard errors associated with commune-level poverty rate estimates

	Mean S. E.	Median S. E.	S. E. ratio	Average # HH	# communes
Phnom Penh	4.0	3.5	35.7	2169	76
Other urban	5.0	4.9	23.8	1345	159
Rural	7.9	7.6	27.4	1289	1359
Cambodia	7.4	7.2	27.4	1337	1594

is presented in the second column (median standard error). The third column (standard error ratio) is the average of the ratio of the standard error to the point estimate (i.e. coefficient of variation). The fourth (average # HH) and fifth columns (# communes) provide the average number of households in the commune and the number of communes in the stratum respectively.

The first three columns provide a general picture of the levels of accuracy. The standard errors are low enough for the results to be useful as proxies, but are high for a number of communes, so policy analysts should take the estimates with caution. At the same time, it should be noted that none of the summary statistics above are perfect. For example, a relatively high level of standard error may not matter if the point estimates are high enough. A commune with the point estimate of the poverty rate of 95 per cent and standard error of 15 per cent is clearly a very poor commune. On the other hand, even if the ratio of the standard error to the point estimates is high, it does not matter when the absolute value of the standard error is low. If the point estimate and standard error were both 0.1 per cent, then the commune is not a poor commune while the ratio is 100 per cent. In practical terms, the size of the commune is also important. Provided that the cost of a programme increases in proportion to the size of the commune, mis-targeting for small communes is a relatively minor issue in terms of efficient use of resources. The statistics above do not incorporate the size of the commune.

One way to address some of these issues is to define a poor commune and non-poor commune by the ratio of the difference between the poverty estimate and a reference level to the standard error. If a poor commune is defined as a commune whose point estimate is higher than the national poverty rate by at least two times the standard error, and if a non-poor commune has the opposite definition, then 48 per cent of all communes can be classified as either poor or non-poor. When a commune cannot be classified, the communes can be aggregated to make the standard error smaller.

Although the magnitudes of the standard errors are not small enough to be ignored, and can be quite high for some communes, the commune-level estimate is accurate enough to make a sharp comparison with the national poverty rate for half of the communes. Even for other communes, the estimates provide useful information for targeting, especially when multiple communes are taken together. It is likely, for example, that net gains from targeting poorer-than-national-average communes are positive. Although the usefulness of the estimates depends upon the purpose to which they are put, given that reliable poverty estimates have previously been produced only at the stratum level, commune-level estimates with this level of accuracy are still very useful. Even when the estimates need to be made at a more aggregated level such as district or even province

to reduce standard errors, the usefulness of the estimates from this exercise will not be undermined as no other reliable estimates are available at this level.¹³

6.4 Extensions

The focus so far has been on poverty measures for the entire population. It is possible to create maps with other measures, such as inequality measures, that can be derived from consumption. It is also possible to derive poverty maps for specific target groups if the census weight for these groups is known. For example, the poverty maps for women and children can be derived by using the number of females and children under the age of five instead of the total household size as the census weight, respectively (see WFP 2002).

7 Applying poverty map to educational programmes

7.1 Education of children and poverty in Cambodia

The poverty map is very useful not only for the identification of the location of the poor but for a variety of other purposes. In this section, we discuss an application of the poverty map to educational programmes. We argue the poverty map is helpful, *inter alia*, for formulating targeting policies for educational programmes such as school feeding-programmes. Education of children is imperative in Cambodia. However, the current situation is not very encouraging. Education policies, as with other policies, have been subject to arbitrary political influences. Ayres (2000: 459) says:

... the provision of formal education in Cambodia has been embraced to build a nation-state that looks modern, yet is concerned almost exclusively with sustaining the key tenets of the traditional polity, where leadership is associated with power and where the nature of the state is perceived to be a function of that power. The result, in terms of education, has been an educational crisis: a significant disparity between the education system and the economic, political, and cultural environments it has been intended to serve.

Educational statistics highlight the abysmal situation in Cambodia. As of 2001, the net enrolment rates (NER) at the primary level for boys and girls are 89 per cent and 83 per cent, respectively. At the secondary level, the statistics are even worse with the NER of 26 per cent for boys and 16 per cent for girls (World Bank 2004). As Bray (1999) argues, Article 68 of Cambodia's constitution, which declares that 'the state shall provide free primary and secondary education to all citizens in public schools', must be taken as a declaration of aspiration rather than reality.

The result of badly planned and implemented education policies, coupled with the lack of resources for education, is the high private cost of public schooling. Bray (1999) studied the private costs of public schooling in Cambodia. Many schools charge fees and other forms of cash payments such as pre-pupil payments demanded for

¹³ Commune, district and provincial level estimates are available in WFP (2002).

construction, repairs and equipment. Teachers often charge their students with fees as salaries are very low. Bray’s study shows that household direct and opportunity costs of schooling are significant factors both for non-enrolment and for drop-outs.

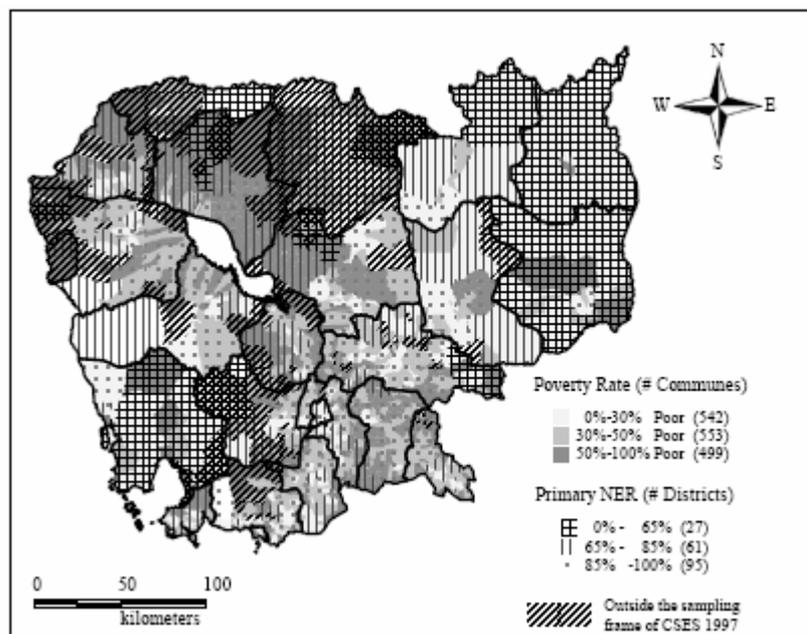
His findings are consistent with the findings of Fujii and Ear (2002). They found children have a better chance of attending the school if their parents are from a wealthier and more educated household. These findings have an important intergenerational implication: children from poor and uneducated households, especially from households with poor and uneducated spouses, are less likely to go to school, which in turn implies that they are more likely to reproduce poor and uneducated children.

To cut this vicious cycle, appropriate education programmes are critical. For example, school feeding-programmes, which supply school children with nutritious food and are designed to provide their parents with incentives for schooling and enhance the nutritional status of children, would be worth consideration. They are also expected to increase the effectiveness in learning, since better concentration in class results from improved nutritional status.

7.2 Finding potential areas for targeting

Education programmes, such as school feeding-programmes, must be appropriately targeted to those areas in which people are poor and the standard of education low. This is particularly the case when the budget for education programmes is severely limited. To identify potential areas for targeting, the poverty map is of great use. We can overlay the education map on the poverty map, and take the intersection between poor areas and low education areas.

Figure 2
Basic education index map overlaid on poverty map



Source: Author

To overlay the educational map on poverty map, we first need to create the education map. The Education Information Management System 2001 dataset was used in order to create the education map. The data were collected annually by the Ministry of Education, Youth and Sport (MoEYS), Department of Planning at school level with support from UNICEF, and contains net enrolment rates at the primary level among other indicators. MoEYS distributes questionnaires to directors of provincial offices of education, youth and sports (PDEYS), and at the provincial level PDEYS provides training on data collection to chiefs of district education office. At the district level, again, chief of district education office arranges training on data collection to directors of all schools within their administrative boundaries. The results are shown in Figure 2. Since we are looking for poor areas with low level of education indicator, we can place priority to the dark and meshed areas. However, it should be emphasized that they are not necessarily the best areas for targeting, because the map only shows the areas most lacking in resources. When targeting school feeding-programmes, other factors should be taken into consideration. Such factors would include the cost of procuring and delivering food, the economies of scale and the capacity of the regional staff.

8 Conclusion

The poverty maps provide invaluable information to policymakers. When there is no reliable information for identifying the poor, targeting policies, if they exist at all, are likely to be inefficient and subject to arbitrary political influences. To deliver assistance to those most in need, policy formulation should be based upon reliable information. The poverty maps presented in this paper give such information at the commune-level. Moreover, the power of poverty maps is multiplied when they are combined with other maps such as the education map using GIS.

As we have discussed in section 7, by superimposing the education map on the poverty map, a policymaker can identify the areas where poverty seems to prevent children from going to school. We have argued that school feeding programmes, for instance, are most likely to be successful if targeted toward such areas, and the method we presented in this study is useful for finding such areas.

The commune-level estimates of poverty rates presented in this paper are reliable enough to be useful. This paper successfully applied the ELL technique to Cambodia and contributed to the empirical evidence of the applicability of the technique. However, it should also be noted that there are errors associated with the estimates and they may be very large for a number of communes. Moreover, the picture depicted here reflects the conditions as of 1998. It should be remembered that analysis of poverty is never static and thus efforts to acquire up-to-date information and monitor changes in poverty will be indispensable for the efficient, effective and timely delivery of assistance. Hence, policymakers should not be misled by the intuitive appeal of poverty maps. The maps presented here can serve as a sound basis to the formulation of targeting policies, but cannot and should not be taken as the sole basis. Other maps and data sources as well as observations from the field should be incorporated in the analysis whenever possible. This is particularly true for the areas outside the sampling frame of CSES 1997.

References

- Alderman, H., M. Babita, G. Demombynes, N. Makhatha, and B. Özler (2002). 'How Small Can You Go? Combining Census and Survey Data for Mapping Poverty in South Africa'. *Journal of African Economies*, 11: 169-200.
- Ayres, D. (2000) 'Tradition, Modernity and the Development of Education in Cambodia'. *Comparative Education Review*, 44 (4): 440-63.
- Baker, J., and M. Grosh (1994). 'Poverty Reduction through Geographic Targeting: How Well Does It Work?'. *World Development*, 22 (7): 983-95.
- Bigman, D., and H. Fofack (2000a). 'Geographical Targeting for Poverty Alleviation: An Introduction to the Special Issue'. *The World Bank Economic Review*, 14 (1): 129-45.
- Bigman, D., and H. Fofack (2000b). 'Geographical Targeting of Poverty Alleviation Programmes: Methodology and Applications in Rural India'. *Journal of Policy Modeling*, 24 (3), 237-55.
- Bigman, D., S. Dercon, D. Guillaume, and M. Lambotte (2000). 'Community Targeting for Poverty Reduction in Burkina Faso'. *The World Bank Economic Review*, 14 (1): 167-94.
- Bray, M. (1999). *The Private Costs of Public Schooling: Household and Community Financing of Primary Education in Cambodia*. Paris: International Institute for Educational Planning and UNICEF.
- Davis, B. (2002). 'Is It Possible to Avoid a Lemon?: Reflections on Choosing a Poverty Mapping Method'. Rome: Agriculture in Economic Development Service, FAO.
- Demombynes, G., C. Elbers, J. Lanjouw, P. Lanjouw, J. Mistiaen, and B. Özler (2002). 'Producing an Improved Geographic Profile of Poverty: Methodology and Evidence from Three Developing Countries'. WIDER Discussion Paper 2002/39. Helsinki: UNU-WIDER.
- Elbers, C., J. O. Lanjouw, and P. Lanjouw (2000). 'Welfare in Villages and Towns'. Timbergen Institute Discussion Paper TI 2000-029/2. Amsterdam: Timbergen Institute.
- Elbers, C., J. O. Lanjouw, and P. Lanjouw (2002). 'Micro-level Estimation of Welfare'. Policy Research Department Working Paper 2911. Washington, DC: World Bank.
- Elbers, C., J. O. Lanjouw, and P. Lanjouw (2003). 'Micro-level Estimation of Poverty and Inequality'. *Econometrica*, 71 (1): 355-64.
- Fujii, T., and S. Ear (2002). 'How Does Spousal Education Matter?: Some Evidence from Cambodia'. *Asian Development Review*, 19 (1): 117-38.
- Hentschel, J., J. O. Lanjouw, P. Lanjouw, and J. Poggi (2000). 'Combining Census and Survey Data to Study Spatial Dimensions of Poverty: A Case Study of Ecuador'. *The World Bank Economic Review*, 14 (1): 147-66.
- Hoddinott, J. (1999). 'Targeting: Principles and Practice'. Technical Guide 9. Washington, DC: International Food Policy Research Institute.
- Kanbur, R. (1987). 'Transfers, Targeting and Poverty'. *Economic Policy*, 4 (1): 112-47.

- Knowles, J. (1998). 'An Updated Poverty Profile for Cambodia, 1997: Technical Report Capacity Development for Socio-Economic Surveys and Planning Project'. Phnom Penh: Ministry of Planning, Royal Government of Cambodia. Mimeo.
- Minot, N. (2000). 'Generating Disaggregated Poverty Maps: An Application to Vietnam'. *World Development*, 28 (2): 319-31.
- Ministry of Planning (MoP) (1998). *A Poverty Profile of Cambodia-1997*. Phnom Penh: Ministry of Planning, Royal Government of Cambodia
- Ministry of Planning (MoP) (2001). *A Poverty Profile of Cambodia-1999*. Phnom Penh: Ministry of Planning, Royal Government of Cambodia.
- National Institute of Statistics (NIS) (1997). 'Selection of Sample Villages Prepared for Cambodia Socio-Economic Survey 1997'. Phnom Penh: NIS, Ministry of Planning, Royal Government of Cambodia. Mimeo.
- National Institute of Statistics (NIS) (1998). *Cambodia Socio-Economic Survey 1997: Technical Report on Survey Design and Implementation*. Phnom Penh: NIS, Ministry of Planning, Royal Government of Cambodia.
- National Institute of Statistics (NIS) (2000a). *Cambodia Socio-Economic Survey 1999: Technical Report on Survey Design and Implementation*. Phnom Penh: NIS, Ministry of Planning, Royal Government of Cambodia
- National Institute of Statistics (NIS) (2000b). *General Population Census of Cambodia 1998: Village Gazetteer*. Phnom Penh: NIS, Ministry of Planning, Royal Government of Cambodia.
- Prescott, N., and M. Pradhan (1997). 'A Poverty Profile of Cambodia'. World Bank Discussion Paper 373. Washington, DC: World Bank.
- Ravallion, M., and K. Chao (1989). 'Targeting Policies for Poverty Alleviation under Imperfect Information: Algorithms and Applications'. *Journal of Policy Modeling*, 11 (2): 213-24.
- Van de Walle, D. (1998). 'Targeting Revisited'. *The World Bank Research Observer*, 13 (2): 231-48.
- World Food Programme (WFP) (2001). *Identifying Poor Areas in Cambodia: Combining Census and Socio-Economic Survey Data to Trace the Spatial Dimensions of Poverty*. Phnom Penh: WFP.
- World Food Programme (WFP) (2002). *Estimation of Poverty Rates at Commune-Level in Cambodia: Using the Small-Area Estimation Technique to Obtain Reliable Estimates*. Phnom Penh: WFP.
- World Bank (2004). *World Development Indicators 2004*. Washington, DC: World Bank.