Micro-Level Estimation of Child Undernutrition Indicators in Cambodia

Tomoki Fujii *

One major limitation to addressing child undernutrition is a lack of the information required to target resources. This article extends the small-area estimation technique of Elbers, Lanjouw, and Lanjouw (2002, 2003) to jointly estimate multiple equations while allowing for individual-specific random errors across equations (in addition to cluster- and household-specific random errors). Estimates of the prevalence of stunting and underweight for children under age 5 in Cambodia from 17 Demographic and Health Survey strata are disaggregated into 1,594 communes by combining the Demographic and Health Survey data. The estimates are consistent with the survey-only estimates at the aggregate and primary sampling unit levels. The accuracy of the commune-level estimates is comparable to the survey-only estimates at the stratum level. The results are robust, and the estimates are useful for policy analysis and formulation. The small-area estimates can be presented in various ways. The strengths of each representation are also discussed. Cambodia, nutrition map, small-area estimation, targeting, undernutrition. JEL classifications: C15, I12, I32, O15

Undernutrition remains a major public health concern in most developing countries. The World Health Organization (WHO 2002) estimates that underweight led to nearly 3.7 million deaths among young children worldwide in 2000. Similarly, Pelletier and others (1994) estimate that about half of child deaths in four developing countries are due to the effects of undernutrition on infectious disease. Undernutrition has also been associated with death and disability later in life, delayed mental development, decreased cognitive and behavioral functioning throughout childhood and adolescence, and poorer performance in school (de Onis, Frongillo, and Blössner 2000; Galler and

*Tomoki Fujii (tfujii@smu.edu.sg) is an assistant professor at the Singapore Management University School of Economics. The author is deeply indebted to Chris Elbers, Jean Olson Lanjouw, and Peter Lanjouw for their advice from the beginning of this research. He also thanks Alain de Janvry, John Donaldson, Livia Montana, Mahadevan Ramachandran, Martin Ravallion, H.E. Kim Saysamalen, Elisabeth Sadoulet, Boreak Sik, H.E. San Sythan, and three anonymous referees. This work was supported by a Singapore Management University research grant (05-C208-SMU-030). Earlier stages of this work were supported by the World Bank and the government of Japan under the Millennium Policy and Human Resource Development Fund.

THE WORLD BANK ECONOMIC REVIEW, VOL. 24, NO. 3, pp. 520–553 doi:10.1093/wber/lhq016
Advance Access Publication December 23, 2010

Barrett 2001; Glewwe, Jacoby, and King 2001; Shariff, Bond, and Johnson 2000).

In Cambodia almost half of children under age 5 are undernourished, as measured by height-for-age or weight-for-age indicators (Cambodian National Institute of Statistics, Cambodian Directorate General for Health, and ORC Macro 2001). Given the grave consequences of undernutrition, the issue deserves serious attention. However, as with many other developing countries, the resources available for improving children's nutrition status are severely limited in Cambodia. Thus, efforts must be made to efficiently allocate the resources available for aid. In particular, geographic targeting—or targeting according to location information—is often easy to administer and implement and can be effective when undernourished children are concentrated in certain locations. However, formulating an effective geographic targeting policy requires knowing the location of undernourished children, information that is often not readily available. The 2000 Cambodia Demographic and Health Survey (CDHS) is representative only at the strata level, which is more aggregated than the province level.[1] Because reasonably reliable estimates of undernutrition are available only at the stratum level or above, the survey estimates are too aggregated to be useful for formulating targeting policies.

This lack of information is a central issue concerning the formulation of targeting policies (Glewwe 1992; Kanbur 1987; Ravallion and Chao 1989). This article aims to overcome the problem by producing commune-level estimates of the prevalence of undernutrition in Cambodia. The estimates can be projected onto maps to allow policymakers to visually identify areas of severe child undernutrition, analyze the current situation of undernutrition, and formulate geographic targeting policies aimed at assisting the neediest individuals more efficiently and transparently. In addition, the method can yield estimates at aggregated levels, which have smaller standard errors than do conventional survey estimates.

To derive the commune-level estimates, the 2000 CDHS data, which include information on child nutrition status but fewer observations, are combined with individual-level Cambodia National Population Census data for 1998, which cover almost all the Cambodian population but lack specific information on child nutrition status. This approach builds on the small-area estimation technique developed by Elbers, Lanjouw, and Lanjouw (2002, 2003). Their methodology is extended to jointly estimate multiple indicators and allow for a richer structure of error terms, a critical step to address issues unique to nutrition indicators.

This article is structured as follows. Section I reviews the measurement and prediction of children's nutrition status. Section II develops the nutrition

1. There are 24 provinces in Cambodia. A province is the most aggregated administrative division after the district, commune, and village. On average, a province has about 89,000 households, a district about 12,000 households, a commune about 1,300 households, and a village about 160 households. Each stratum has an average of 125,000 households.

mapping methodology. Section III discusses the data. Section IV presents the results. And section V analyzes the results and compares them with previous estimates.

I. Measuring and Predicting Child Nutrition Status

To noninvasively and inexpensively measure undernutrition, researchers—including economists, nutritionists, and epidemiologists—use anthropometry. Among the most commonly used anthropometric measures are weight-for-height, weight-for-age, and height-for-age *Z*-scores, which measure the number of standard deviations between an individual's value of the anthropometric indicator and the median of the National Center for Health Statistics growth reference population of the same sex and age or height group. Deficiencies in weight-for-height, weight-for-age, and height-for-age *Z*-scores are respectively called "wasting," "underweight," and "stunting." The conventional cutoff of −2 is used to calculate the prevalence of undernutrition. For example, the prevalence of stunting is defined as the number of children with a height-for-age *Z*-score below −2, divided by the total number of children (see Dibley and others 1987a, b; Waterlow and others 1997; WHO Working Group 1986, 1996; and Sahn and Stifel 2003 for further discussion on *Z*-scores).

As WHO Working Group (1986) points out, there are several obvious differences among these measures. First, it is possible to lose weight but not height. Second, linear growth is slower than growth in body mass. Third, catch-up in height is possible but takes a long time, even in a favorable environment. Thus, wasting reflects acute, or short-term, undernutrition, whereas stunting reflects chronic, or long-term, undernutrition, with underweight somewhere in between. Given these differences, it should not be surprising when patterns of wasting and stunting are different. In fact, Victora (1992) finds no systematic pattern that holds for an international population between stunting and wasting.

Although the measurement and cutoff used to define undernutrition are widely accepted in economics, public health, and nutrition studies, the definition does not distinguish the true diseases of undernutrition from normal shortness or thinness. However, as suggested by Wright, Ashenburg, and Whitaker (1994), the undernutrition measures used here can be interpreted as the risk of the adverse effects of undernutrition. The methodology does not depend on the measurement or cutoff, so other measurements or cutoffs can be used.

Following Pradhan, Sahn, and Younger (2003), a standardized height and weight are used, defined as the *Z*-scores converted back to the corresponding height and weight of the reference age-sex group of 24-month-old girls. Standardization preserves all the desirable properties of the original *Z*-scores, and standardized heights and weights are always positive for practically possible values of *Z*-scores, so more meaningful measures can be computed (Foster, Greer, and Thorbecke 1984).[2] That is, the standardized height and weight corresponding to a −2 Z-score is analogous to the poverty line Z, a cutoff below which an individual is considered poor. The Foster, Greer, and Thorbecke (1984) measure of undernutrition can be defined by simply replacing consumption with the standardized height or weight.

As noted by Pradhan, Sahn, and Younger (2003), the choice of the age-sex group is arbitrary. Therefore, 6-month-old girls, 60-month-old girls, and 24-month-old boys are chosen as the reference age-sex groups. For both stunting and underweight the pairwise correlation of the estimated commune-level prevalence of undernutrition for any two reference groups is at least 99.99 percent. Thus, as with Pradhan, Sahn, and Younger (2003), the choice of the reference age-sex group does not substantively change the results.

Now, let $y_i^{(1)}$ and $y_i^{(2)}$ be individual i's standardized height and weight, and $z^{(1)}$ and $z^{(2)}$ be the standardized height and weight corresponding to a −2 Z-score. The Foster, Greer, and Thorbecke (1984) measure of undernutrition with parameter θ can be written as follows:

$$(1) \qquad P^{\theta,(k)} = \sum_i \mathrm{Ind}(y_i^{(k)} < z^{(k)}) \cdot \left(\frac{z^{(k)} - y_i^{(k)}}{z^{(k)}} \right)^{\theta} \qquad (k = 1, 2)$$

where Ind(·) denotes the indicator function, $P^{0,(1)}$ and $P^{0,(2)}$ are the prevalence of stunting and underweight, and $P^{1,(1)}$ and $P^{1,(2)}$ are the undernutrition gap for stunting and underweight.

Because the methodology is applicable to cases with more than two undernutrition indicators, current nutrition status as measured by weight-for-height could be included in the analysis in principle. But it is not included because a regression model of weight-for-height with sufficient explanatory power could not be constructed, a problem not limited to Cambodia. Building a good model of weight-for-height is often difficult because the individual or household characteristics observed in a typical Demographic and Health Survey—such as education, demographic composition, and housing conditions, which do not vary much in the short run—cannot capture fluctuations in the short-term nutrition status of children. Using data from Morocco, Pakistan, South Africa, and Vietnam, Alderman (2000) found that the explanatory power of the regression models among all the anthropometric indicators discussed above was smallest for weight-for-height.

Because the methodology is built on the association between anthropometric indicators and other socioeconomic and geographic indicators, this section briefly reviews previous studies that describe the relationship between anthropometric indicators and other indicators. Various studies—including Frongillo,

2. Another purpose behind this decision is to calculate health inequality measures, which are discussed in Fujii (2007a). The reference group is the same as that in Pradhan, Sahn, and Younger (2003).

de Onis, and Hanson (1997), Haughton and Haughton (1997), Haddad and others (2003), and Li and others (1999)—indicate the potential importance of the age and sex of the child, female education, and access to potable water in explaining the nutrition status of children. Therefore, these variables are also included in the pool of potential explanatory variables. Other important variables, such as the anthropometry of the mother and household income, are not included because the census data lack these variables.

In addition to the individual- and household-level characteristics mentioned above, various studies suggest the importance of including variables that characterize the community or the location. Alderman (2000) and Curtis and Hossain (1998) show that the location of residence may be important in explaining anthropometric indicators, even after controlling for some observable characteristics. The model includes several village- and commune-level variables in the pool of potential explanatory variables.

Finally, Zeini and Casterline (2002) suggest that conditional correlations may exist at various levels. Using the 2000 Egypt Demographic and Health Survey, they find such a correlation at the individual level (that is, across various anthropometric indicators), at the household level (that is, across siblings), and at the location level (that is, within the same neighborhood) for some regions. The anthropometric model used here also allows for such conditional correlations, as elaborated in the section II.

II. Methodology

The methodology is similar to the small-area estimation by Elbers, Lanjouw, and Lanjouw (2002, 2003), which has been used to analyze poverty and inequality. Both methodologies combine survey data with unit-record census data to obtain estimates at a lower level of aggregation than the survey permits. Elbers, Lanjouw, and Lanjouw's (2002, 2003) small-area estimation approach was first applied to data from Ecuador (Hentschel and others 2000) and has subsequently been applied to several other countries (see, for example, Alderman and others 2002; Demombynes and Ozler 2005; and Elbers and others 2007). Applications to Cambodia can be found in Fujii (2006, 2007b, 2008).

The basic idea is straightforward. First, the parameters of the anthropometric models are estimated using the survey dataset. The estimates are then used to impute the anthropometric indicators to each census record. The imputed anthropometric indicators are then aggregated to arrive at estimates of undernutrition measures for small areas. An important feature of the current study and Elbers, Lanjouw, and Lanjouw (2002, 2003) is the explicit treatment of the standard errors of these aggregate estimates. To account for idiosyncratic error and model error, the disturbance terms and the model coefficients are repeatedly drawn in a Monte Carlo simulation, the details of which are discussed later in this section.

Four differences between Elbers, Lanjouw, and Lanjouw's (2002, 2003) study and the current study warrant discussion here. First, the type of the survey dataset used for the estimation differs. Elbers, Lanjouw, and Lanjouw (2002, 2003) focus on consumption or income data from a socioeconomic survey, whereas the current study uses anthropometric measures from a Demographic and Health Survey. Second, the unit of analysis differs. Consumption data are usually produced at the household level, whereas anthropometric measures are produced at the individual level. Under the Elbers, Lanjouw, and Lanjouw (2002, 2003) approach, disturbance terms are decomposed into cluster- and household-specific effects. The study allows for an unobserved individual-specific effect in addition to cluster- and household-specific effects. Third, unlike Elbers, Lanjouw, and Lanjouw (2002, 2003), the current study makes finite-sample corrections. The number of children under age 5 in a household is limited, with no more than two for most households in Cambodia. This means that large-sample properties cannot be relied on when estimating the individual-specific effect, so finite-sample correction is crucial. Fourth, Elbers, Lanjouw, and Lanjouw (2002, 2003) consider only one equation, whereas the current study simultaneously estimates multiple equations. The current study also allows for the correlation of individual-specific effect across indicators, referred to as the intrapersonal correlation. This correlation must be taken into account to reproduce the correlation across indicators at aggregate levels.

The following subsections describe the key assumptions of the methodology and discuss the main ideas behind the parameter estimation and simulation. The technical details are in the appendix.

Key Assumptions

As with the Elbers, Lanjouw, and Lanjouw (2002, 2003) estimation, the current study makes a few important assumptions. First, the predictors, or the explanatory variables used to impute the anthropometric indicators, are assumed to be measured in the same way in the census as in the survey. For example, if a household is included in both the census and the survey, the predictors for this household must be equal in the two samples. Unfortunately, the equality of predictors in the two samples cannot be tested at the household level because the data are anonymized.

When the underlying populations for the two samples are identical, whether the predictors have the same distribution in the two samples can be tested. In practice, the two underlying populations are often not identical because the census and the survey data are collected at different points in time. As a result, it usually must be assumed that the same model is applicable to the two underlying populations. Although the comparability of measurement cannot be formally tested in this case, it is still useful to look at the distributions of the predictors. For example, variables can be identified that have very different distributions in the two samples. If the difference for a certain variable is too great to be explained solely by the timing of data collection, the variable can be deleted from the list of candidate predictors.

Second, it is assumed that no correlation exists between the random effects specific to the cluster (location). This is an important assumption because intercluster correlations lead to the underestimation of the standard errors, as pointed out by Tarozzi and Deaton (2009), who also show that the underestimation may be substantial under plausible conditions. However, some intercluster correlations may be captured by the fixed effects at a more aggregated level, as is often done in practice.

Further, some empirical evidence indicates the appropriateness of the Elbers, Lanjouw, and Lanjouw (2002, 2003) estimator. Demombynes and others (2007) test the Elbers, Lanjouw, and Lanjouw (2002, 2003) estimates by comparing them with actual observations in Mexico and find that when the model includes location variables to keep unobserved location effects small, the Elbers, Lanjouw, and Lanjouw (2002, 2003) approach yields appropriate standard errors. However, because the small-area estimates are based on a target population that consists of randomly selected (noncontiguous) villages, the intercluster correlation in Demomybnes and others (2007) is likely to be small by construction. Elbers, Lanjouw, and Leite (2008) overcome this by using data that include contiguous locations in the state of Minas Gerais, Brazil. Their results also indicate that the Elbers, Lanjouw, and Lanjouw (2002, 2003) standard errors are appropriate.

Because of lack of data, the current study's method cannot be tested in the same way that Demomybnes and others (2007) and Elbers, Lanjouw, and Leite (2008) can be, but the robustness of the results can be checked by changing the definition of clusters, as discussed later. In the baseline simulation a cluster is taken to be a village, the primary sampling unit for the survey. The robustness of the results is checked by letting the cluster be a more aggregated administrative unit such as a commune or a district.

Estimation

To describe the important issues in estimation and simulation, some notations must first be introduced. The cluster, household, and individual are denoted by the subscripts c, h, and i respectively. $y_{chi}^{(k)}$ is the kth $(1 \leq k \leq K)$ anthropometric indicator of interest, and $\mathbf{x}_{chi}^{(k)}$ is a $d^{(k)}$ vector of observable characteristics that are used as a predictor of $y_{chi}^{(k)}$. In the empirical application $K = 2$, with $k = 1$ and $k = 2$ as the standardized height and weight, respectively. The following linear approximation to the conditional distribution of $y_{chi}^{(k)}$ is considered.

$$\begin{aligned} y_{chi}^{(k)} &= E\left[y_{chi}^{(k)} | \mathbf{x}_{ch}^{(k)}\right] + u_{chi}^{(k)} = \left[\mathbf{x}_{chi}^{(k)}\right]^T \beta^{(k)} + u_{chi}^{(k)} \\ &= \left[\mathbf{x}_{chi}^{(k)}\right]^T \beta^{(k)} + \eta_c^{(k)} + \varepsilon_{ch}^{(k)} + \delta_{chi}^{(k)}, \end{aligned} \tag{2}$$

where $\beta^{(k)}$ is a $d^{(k)}$ vector of parameters and $u_{chi}^{(k)}$ is a disturbance term that consists of the cluster-specific effects $\eta_c^{(k)}$, the household-specific effects $\epsilon_{ch}^{(k)}$, and the individual-specific effects $\delta_{chi}^{(k)}$. These three components of the disturbance term are assumed to be uncorrelated with each other and across clusters, households, and individuals. The cluster- and individual-specific effects are assumed to be homoskedastic, but heteroskedasticity of the household-specific effect is allowed for. Finally, the cluster- and household-specific effects are assumed to be uncorrelated across indicators, but the correlation of individual-specific effect across indicators (intrapersonal correlation) is allowed for.

This specification is motivated in part by the studies discussed in the previous section. Introspection also supports this specification. First, including the cluster-specific effect is important because some important and unobservable determinants of undernutrition, such as the availability of clean water and the risk of infectious diseases like malaria, are cluster specific. Second, the inclusion of household-specific effects may also be important. Unobserved household characteristics, such as the adequacy of child care in the household and parental genetic information, may affect the observed nutrition status of children. Third, because unobservable heterogeneity within the household may affect both indicators, it is important to allow for intrapersonal correlation. For example, the distribution of food within the household may differ across siblings, but such heterogeneity is unobservable. In contrast, both height and weight would be affected by food distribution.

The specification is also driven by data limitations. In principle, both the cluster- and individual-specific effects may be heteroskedastic. However, cluster-level heteroskedasticity is difficult to estimate because the number of clusters in the survey is limited. Once the heteroskedasticity of the household-specific effect is allowed for, it is difficult to identify the heteroskedasticity of the individual-specific effect because there are typically only one or two children under age 5 in each household. Further, the intrapersonal correlation is also difficult to estimate accurately when the individual-specific effect is heteroskedastic.

To estimate the model coefficient and the distribution of the disturbance term, a (feasible) generalized least squares estimation is conducted. In the first-stage regression an ordinary least squares regression of $y_{chi}^{(k)}$ on $x_{chi}^{(k)}$ was run. Then, the ordinary least squares residual was used to estimate the variance of each component of the error as well as the correlation of the individual effect across indicators.

Because of differences in the structure of the disturbance term and the nature of the data, the strategy to estimate the variance-covariance matrix of u is markedly different from that used by Elbers, Lanjouw, and Lanjouw (2002, 2003). They assume a large number of households in each cluster, so the cluster-specific effect is estimated at the mean of the ordinary least squares residual within each cluster, and the household-specific effect is estimated at the deviation of the ordinary least squares residual from the cluster mean.

However, the same method cannot be applied to separate the individual-specific effect from the household-specific effect in the current study because the number of children under age 5 in a household is small. A finite-sample correction must thus be employed. To this end, the deviation of the disturbance term from the household mean is considered for each individual. The expected value of this deviation squared and then multiplied by the inverse of the finite-sample correction factor (that is, 1 minus the reciprocal of the number of children under age 5 in the household) equals the variance of the individual-specific effect. Thus, the variance of the individual-specific effect can be estimated by replacing the disturbance term with the ordinary least squares residual (see equation [A1] in the appendix). Similarly, the product of the deviation for different anthropometric indicators is used to estimate the covariance of the individual-specific effect (see equation [A2] in the appendix). For the variance of the cluster-specific effect, the difference between the average of the squared cluster mean of u and the average of the squared household mean of u is used (equation [A3] in the appendix).

Because some survey clusters include only a few households, it is difficult to separate the household-specific effect from the cluster-specific effect. Thus, to obtain the variance of the household-specific effect, the sum of the household- and cluster-specific effects is examined first (equation [A4] in the appendix). This sum is heteroskedastic, and the only source of heteroskedasticity is the household-specific effect. Therefore, a heteroskedastic regression for the sum is run first; then, the variance of the household-specific effect is found by subtracting the variance of the cluster-specific effect. As with Elbers, Lanjouw, and Lanjouw (2002, 2003), a logistic-type transformation (equation [A5] in the appendix) is used to ensure that the predicted variance of the household-effect is bounded and non-negative, and the variance of the household-specific effect for each household (equation [A6] in the appendix) is estimated.

With all components of the variance-covariance matrix of the disturbance term, the generalized least squares regression can be run. The distribution of each component of the disturbance term from the empirical distribution of the ordinary least squares residual is also estimated. With all these estimation results, the simulation calculations can proceed.

Simulation

Using the estimation results, the anthropometric indicators are imputed for each child under age 5 in the census. The imputed value is subject to two sources of error: the model error, which arises from the error in the estimation of model coefficients, and the idiosyncratic error, which arises from the fact that even if the true β is known, the imputed value would not equal the actual anthropometric indicator because of the nonsystematic component u in equation (2).

As with Elbers, Lanjouw, and Lanjouw (2002, 2003), these two types of errors are taken into account through Monte Carlo simulation. In each round of simulation the coefficient $\tilde{\beta}^{(k)}$ for all k is simultaneously drawn and then used to impute the systematic component in equation (2). The individual-, household-, and cluster-specific effects for each individual, household, and cluster are then drawn from their estimated distributions, where the draw is done simultaneously for all indicators from their standardized empirical distributions. The variance of the household-specific effect is imputed with the heteroskedastic model. With all these parameters, anthropometric indicators have been imputed for each census record in each round of the simulation (see equation [A7] in the appendix).

The imputed anthropometric indicators can now be used to obtain estimates of the prevalence of undernutrition for small geographic areas (such as communes) in each round of the simulation. Taking the mean and standard deviation of these estimates over all rounds of the simulation yields the point estimates and the standard errors for small geographic areas.

III. Data

The basic building blocks for this study comprise a survey dataset, a census dataset, and a dataset of geographic variables. The survey dataset used is CDHS 2000, which was designed to collect demographic and health information on the Cambodian population, with a particular focus on women of childbearing age and young children. The sample covered 12,236 households in 17 strata across the country. Data collection took place between February and July 2000. The 2000 CDHS has a three-stage sampling design, where the primary sampling unit is a village and the secondary sampling unit is a segment (a block of about 10 households) or a collection of segments. In the dataset used for this study, the village to which each household belongs can be identified but the segment cannot.

In addition to detailed information about each household, its members, and its housing characteristics, a quarter of households were systematically selected to participate in the anthropometric data collection. All children under age 60 months in the subsampled households were weighed and measured. After excluding children for whom information on height or weight was missing or implausible, 3,596 observations were used for this analysis (for further details, see Cambodian National Institute of Statistics, Cambodian Directorate General for Health, and ORC Macro 2001). The height-for-age and weight-for-age Z-scores were derived first and then converted to the height and weight of 24 month-old girls with the same Z-scores.

The second source of data is the Cambodian National Population Census, the first population census to be conducted in Cambodia since 1962. The census covered virtually all people in Cambodia at the reference time of midnight of March 3, 1998.[3] The census data contain information on housing

3. Due to military operations, about 0.5 percent of the population was not covered.

characteristics and information on each usual household member and visitors present on the reference night, including the relationship to the head of household, sex, age, marital status, migration status, literacy, education level, and employment status. The census dataset contains more than 1.4 million records of children under age 5.

The two samples have similar distributions (table 1). There is no clear evidence that the assumption of comparability of measurement is violated. The standards of living in Cambodia are low. Housing conditions are generally poor. Most households do not have electricity or toilet facilities, and their houses are built with poor materials. The educational attainment of household heads is also low.

In addition to the survey and the census, a set of geographic indicators is used in this analysis. Because Cambodia has a rich collection of geographic data, indicators could be generated on a range of characteristics, including distance calculations, land use and land cover information, climate indicators, vegetation, agricultural production, and flooding. After compiling numerous datasets from various sources into a geographic information system, these indicators were generated for all villages and communes in Cambodia. Coarse-resolution data were summarized at the commune level, and high-resolution data at the village level. The distance from the center of villages to roads, other towns, health facilities, and major rivers was calcualted. Indicators based on satellite data with varying temporal resolutions included land use within the commune (agricultural, urban, forested, and so forth), the normalized difference vegetation index to proxy agricultural productivity, and the degree to which the area was lit by nighttime lights as a proxy of urbanization. Relatively stable indicators including soil quality, elevation, and various 30-year average climate variables were derived from other composite datasets. Village-level means were also generated from the census data. These means do not have to be taken from the variables that also exist in the CDHS dataset because the village-level means, as with other geographic variables, can be linked to both the census and the survey datasets. Including these geographic variables and their cross terms with other individual- and household-level variables has substantially improved the ability to explain the variation in anthropometric indicators.

IV. Results

An anthropometric model was constructed in each of the five zones ("ecozones") of Cambodia: Coastal, Plain, Plateau, Tonle Sap, and Urban. All are rural except the Urban ecozone. Provinces with similar agroclimatic and sociocultural characteristics were combined because some of the strata had too few observations to carry out a meaningful analysis.[4] Regressions were run

4. This uses the same definition of the ecozones as WFP (2001).

TABLE 1. Key Summary Statistics for the Survey and Census

	Survey		Census	
Variable	Mean	Standard deviation	Mean	Standard deviation
Child's demographics				
Male	0.51	(0.50)	0.51	(0.50)
Age = 0	0.20	(0.40)	0.16	(0.36)
Age = 1	0.17	(0.37)	0.19	(0.39)
Age = 2	0.19	(0.39)	0.19	(0.39)
Age = 3	0.22	(0.41)	0.23	(0.42)
Age = 4	0.22	(0.42)	0.23	(0.42)
Household demographics				
Age of household head	38.54	(11.78)	36.40	(11.19)
Age of spouse	30.83	(13.91)	28.23	(14.45)
Household size	6.26	(2.17)	6.08	(2.30)
Male household head	0.84	(0.37)	0.82	(0.38)
Main cooking fuel				
Charcoal	0.06	(0.23)	0.04	(0.20)
Firewood	0.92	(0.27)	0.93	(0.26)
Other	0.02	(0.15)	0.03	(0.17)
Drinking water source				
Piped water	0.04	(0.19)	0.04	(0.20)
Tube or piped well	0.20	(0.40)	0.14	(0.35)
Dug well	0.36	(0.48)	0.42	(0.49)
Surface water	0.32	(0.47)	0.30	(0.46)
Bought	0.04	(0.19)	0.07	(0.26)
Other	0.03	(0.18)	0.03	(0.16)
Roof material				
Wood or plastic	0.48	(0.50)	0.52	(0.50)
Metal	0.24	(0.43)	0.24	(0.43)
Rock or other	0.28	(0.45)	0.24	(0.43)
Floor material				
Dirt (including earth, sand, and clay)	0.10	(0.30)	0.16	(0.37)
Wood (including bamboo)	0.83	(0.37)	0.72	(0.45)
Other materials	0.06	(0.24)	0.12	(0.32)
Other housing conditions				
Have electricity	0.13	(0.33)	0.12	(0.33)
Have toilet facility	0.15	(0.36)	0.11	(0.32)
Household head's education				
None	0.24	(0.43)	0.29	(0.45)
Incomplete primary	0.44	(0.50)	0.40	(0.49)
Complete primary	0.07	(0.26)	0.06	(0.24)
Incomplete secondary	0.22	(0.42)	0.14	(0.35)
Complete secondary	0.02	(0.14)	0.05	(0.22)
College or higher	0.01	(0.08)	0.06	(0.23)
Number of observations	3,596		1,453,286	

Source: Author's calculation based on Cambodian Demographic and Health Survey and census data.

separately for each ecozone, using the regressors that have similar marginal distributions within the ecozone. Although individual- and household-level variables explain only 20–30 percent of the variations in the standardized height and weight, the explanatory power was increased to over 40 percent when geographic variables, interaction terms, and other transformations of variables were included in the set of potential explanatory variables. To avoid overfitting the data, the model was as parsimonious as possible. The robustness of the regression coefficients was checked by randomly dropping some households or clusters, as was done in Elbers, Lanjouw, and Lanjouw (2002, 2003).

The first-stage ordinary least squares regression results for standardized height and weight for the Coastal ecozone illustrate the importance of geographic variables and interaction terms. Although causal interpretations should not be made, the coefficients on the age dummy variables are all negative, which means that a zero-year-old child is healthier than an older child after controlling for some other factors (tables 2 and 3). This does not reflect a dramatic improvement in nutrition status in 2000, but it does reflect that infants are less likely to be exposed to contaminated food before they are weaned, so they are less likely to suffer from diarrhea.

The point estimate of the variance of the cluster-specific effect was 0 in all strata. However, because the ordinary least squares residuals are bootstrapped, $(\tilde{\sigma}^{(k)}_{\eta,(r)})^2$ was strictly positive in some rounds of the simulation. The average proportion of the individual-specific effect to the variance of the disturbance term as a whole was high. The ratio of the variance of the individual-specific effect to the overall residuals, $\frac{(\hat{\sigma}^{(k)}_{\delta})^2}{(\hat{\sigma}^{(k)}_{u})^2}$, was over 0.86 for standardized height and over 0.69 for standardized weight. This means that the individual-specific effect dominates the household- and cluster-specific effects. In all ecozones the magnitude of intrapersonal correlations was high, ranging from 0.42 to 0.53. Therefore, the magnitude of the intrapersonal correlation is substantial.

After the predictions for standardized height and weight for each child in the census were made in each round of simulation, they were aggregated to arrive at the estimate of the prevalence of stunting and underweight at the commune level. Because of missing data in the census and geographic datasets for a small number of communes, estimates were obtained for 1,594 of the 1,616 communes in Cambodia.

Table 4 reports the prevalence of stunting and underweight estimated with the CDHS only and the corresponding small-area estimates[5] as well as the quartiles and maximum of the standard errors for commune- and district-level

5. Clustering at the segment level is ignored for three reasons. First, segments are not observed in the CDHS or census, so they cannot be modeled. Second, spatial heterogeneity within a village is generally limited in Cambodia. Third, even if segment-level clustering exists, it is at least partly captured in the cluster-specific effect in the model.

TABLE 2. First-stage Ordinary Least Squares Regression Results for Standardized Height

Variable	Full model Estimate	Full model Standard error	Model without interaction terms Estimate	Model without interaction terms Standard error	Model without interaction terms and geographic variables Estimate	Model without interaction terms and geographic variables Standard error
Intercept	-297.1	(103.8)	-296.3	(104.6)	83.63	(1.384)
Maximum household head education in years	0.388	(0.148)	0.561	(0.132)	0.717	(0.134)
Non-dirt/wood floor	-3.994	(1.817)	-4.282	(1.819)	-3.453	(1.876)
Piped water	-2.336	(1.506)	-2.680	(1.528)	-2.398	(1.574)
Dug well	-3.960	(0.953)	-1.794	(0.712)	-1.373	(0.739)
Male child	-1.469	(0.633)	-1.469	(0.635)	-1.283	(0.660)
Ratio of educated members in household	-0.037	(0.018)	-0.035	(0.018)	-0.038	(0.019)
Number of girls under age 5 in household	5.030	(1.907)	-3.968	(1.133)	-2.681	(1.156)
One-year-old child	-3.794	(1.038)	-3.710	(1.061)	-3.733	(1.102)
Two-year-old child	-2.899	(1.024)	-2.674	(1.038)	-2.514	(1.084)
Three-year-old child	-10.18	(1.183)	-8.103	(0.904)	-8.512	(0.931)
Four-year-old child	-11.41	(2.005)	-6.268	(0.898)	-6.479	(0.934)
Ratio of households in village with a college educated member	415.6	(104.2)	413.4	(104.9)		
Ratio of households in village with a college educated spouse	-32.38	(8.288)	-33.25	(8.427)		
Ratio of household heads with irregular employment in village	552.7	(279.5)	518.2	(285.4)		
(Maximum household education in years)*(Four-year-old child)	0.537	(0.254)				
(Number of girls under age 5 in household)*(Ratio of households in village with a college educated member)	-5.037	(1.908)				
(Male child)*(Three-year-old child)	3.703	(1.560)				
(Male child)*(Four-year-old child)	3.839	(1.643)				
(Male child)*(Number of girls under age 5 in household)	0.006	(0.002)				
F	11.09		12.87		12.31	
R^2	0.406		0.365		0.300	

Note: Number of observations is 326. All regressions were run with the sample weight.

Source: Author's analysis based on the Cambodia Demographic and Health Survey for 2000 and the Cambodia National Population Survey for 2000.

Table 3. First-Stage Ordinary Least Squares Regression Results for Standardized Weight

	Full model		Model without interaction terms		Model without interaction terms and geographic variables	
Variable	Estimate	Standard error	Estimate	Standard error	Estimate	Standard error
Intercept	-90.26	(19.60)	-86.95	(20.35)	11.87	(0.29)
Maximum female education in years	0.134	(0.047)	0.038	(0.026)	0.061	(0.026)
Dug well	-3.709	(1.672)	-0.230	(0.131)	-0.299	(0.135)
Charcoal for cooking	0.648	(0.263)	0.464	(0.270)	0.181	(0.277)
Wooden floor	1.272	(0.269)	0.515	(0.158)	0.274	(0.163)
Number of children between ages 1 and 4	0.244	(0.124)	0.277	(0.128)	0.139	(0.136)
Household head's age	-0.021	(0.005)	-0.020	(0.005)	-0.018	(0.005)
One-year-old child	-2.386	(0.330)	-1.640	(0.227)	-1.597	(0.243)
Two-year-old child	-1.198	(0.198)	-1.205	(0.206)	-1.196	(0.224)
Three-year-old child	-1.786	(0.186)	-1.810	(0.193)	-1.951	(0.205)
Four-year-old child	-1.641	(0.190)	-1.675	(0.197)	-1.633	(0.213)
Change in agricultural land cover in commune between 1993 and 1997	-0.057	(0.017)	-0.042	(0.017)		
Ratio of households in village with a college educated member	107.6	(19.74)	102.7	(20.47)		
Ratio of professionals in village	-7.355	(3.350)	-5.771	(3.466)		
Ratio of household heads with irregular employment in village	7.673	(3.148)	9.549	(3.102)		
Flood prone area in 1996	0.645	(0.242)	0.535	(0.246)		
Ratio of married household heads in village	-7.157	(1.480)				
(Dug well)*(Ratio of married household heads in village)	4.172	(1.996)				
(Maximum female education in years)*(Wooden floor)	-0.184	(0.053)				
(Maximum female education in years)*(One-year-old child)	0.225	(0.070)				
F	14.65		14.52		14.18	
R^2	0.475		0.428		0.309	

Note: Number of observations is 326. All regressions were run with the sample weight.

Source: Author's analysis based on the Cambodia Demographic and Health Survey for 2000 and the Cambodia National Population Survey for 2000.

Table 4. Ecozone-level Comparison between the Survey-Only and Small-Area Estimation Estimates with Various Levels of Clustering

	$P^{0,(1)}$ (% stunted)				$P^{0,(2)}$ (% underweight)			
		Small area estimation				Small area estimation		
Ecozone	Survey only Village Estimate (standard error)	Village Estimate (standard error)	Commune Estimate (standard error)	District Estimate (standard error)	Survey only Village Estimate (standard error)	Village Estimate (standard error)	Commune Estimate (standard error)	District Estimate (standard error)
Coastal	47.21 (5.05)	47.05 (3.40)	47.04 (3.04)	47.50 (2.82)	38.95 (5.33)	39.03 (3.32)	39.23 (3.26)	39.34 (3.00)
Plain	47.58 (2.75)	50.81 (1.78)	50.80 (1.78)	50.43 (1.60)	47.80 (2.71)	46.57 (1.70)	45.93 (1.71)	46.41 (1.46)
Plateau	47.10 (3.08)	46.86 (1.39)	46.49 (1.38)	45.30 (1.32)	46.37 (3.16)	45.92 (1.59)	45.17 (1.61)	45.84 (1.56)
Tonle Sap	42.87 (2.38)	43.95 (2.09)	43.86 (1.80)	45.05 (1.58)	45.84 (2.46)	43.12 (1.93)	42.47 (1.92)	43.64 (1.58)
Urban	37.89 (2.92)	37.03 (1.56)	39.24 (1.72)	37.92 (1.68)	39.58 (2.65)	38.69 (1.60)	41.24 (1.58)	39.56 (1.61)
First quarter of standard errors	—/—	2.95/2.19	3.22/2.19	3.11/2.34	—/—	3.03/2.24	3.28/2.24	3.05/2.41
Second quarter of standard errors	—/—	3.55/2.58	3.90/2.59	3.74/2.96	—/—	3.54/2.57	3.79/2.61	3.76/2.95
Third quarter of standard errors	—/—	4.46/3.13	4.82/3.04	4.71/3.59	—/—	4.48/3.17	4.70/3.03	4.68/3.69
Maximum of standard errors	—/—	21.59/8.43	17.59/8.19	19.54/7.04	—/—	15.63/6.57	15.63/6.07	16.71/7.59

Note: The standard errors for the survey-only estimates were calculated by a 100-time two-stage bootstrapping simulation.

Source: Author's analysis based on the Cambodia Demographic and Health Survey for 2000 and the Cambodia National Population Census for 1998.

estimates. For example, the median standard error of commune- and district-level estimates of the prevalence of stunting in the baseline estimate is 3.55 percent and 2.58 percent.

Five remarks about table 4 are in order. First, the differences between the survey-only estimates and the small-area estimates are within two standard errors of the survey-only estimates in all cases, so the differences in the ecozone-level estimates can be attributed to random errors. The survey-only estimates of the prevalence of undernutrition were compared with the baseline small-area estimates for 24 provinces.[6] The difference between the survey-only estimates and the small-area estimates is within one standard error of the survey-only estimates in most cases (table 5). Further, the survey-only estimates are not significantly different from the small-area estimates in 47 of the 48 comparisons at a 5 percent significance level. There is no obvious explanation for the exception—the prevalence of stunting in Banteay Mean Chey province. But given the number of comparisons, this exception can be attributed to random error.

Second, the commune-level estimates have reasonably small standard errors, with a median standard error of less than 4 percent. The magnitudes of the standard errors for the commune-level estimates are comparable to the survey-only estimates. However, in all cases in table 4, some communes have estimated standard errors over 15 percent. Thus, there are some communes for which estimates are poor, even though reasonably accurate estimates were obtained for a majority of communes.

Third, even if the estimates have very high standard errors, they are still useful. The idiosyncratic error tends to decrease when there are more communes in a target group. So, the standard errors for the target group as a whole can be much lower than the standard error for each commune in the target group, especially when there are many communes in the target group. Therefore, if a proposed nutrition intervention delivers assistance to a large number of communes, high standard errors for the individual communes are not necessarily worrisome.

Fourth, the level at which clustering occurs has no systematic impact on the ecozone-level estimates of the prevalence of undernutrition. The point estimates are not affected because clustering changes only the error structure and thus has no systematic impacts on the estimated model coefficients. The effects of varying the clustering unit on the standard errors are ambiguous. If the correlation of the disturbance term within a cluster is fixed, the standard error tends to be higher when the clustering unit is more aggregated. Thus, the estimates based on village-level clustering may be too optimistic. But the observed correlation of disturbance terms within a cluster tends to be smaller when a more

6. Even though the survey is not representative at the province level, the estimates are reported at this level because the province is the administrative unit most relevant to policymaking.

TABLE 5. Provincial-Level Comparison of the Prevalence of Stunted and Underweight Children

Province	$P^{0,(1)}$ (% stunted) Survey only Estimate (standard error)	$P^{0,(1)}$ (% stunted) Small-area estimation Estimate (standard error)	$P^{0,(2)}$ (% underweight) Survey only Estimate (standard error)	$P^{0,(2)}$ (% underweight) Small-area estimation Estimate (standard error)
Banteay Mean Chey	28.86	41.75	39.48	45.86
	(4.22)	(2.02)	(4.99)	(2.16)
Battambang	36.28	38.82	37.17	38.23
	(5.42)	(2.18)	(4.99)	(2.26)
Kampong Cham	48.20	52.00	47.80	46.62
	(5.58)	(1.93)	(5.41)	(2.03)
Kampong Chhnang	45.94	47.85	46.12	44.05
	(4.10)	(2.09)	(4.69)	(1.85)
Kampong Speu	44.50	41.66	44.04	44.23
	(4.45)	(1.87)	(4.21)	(2.05)
Kampong Thom	47.19	43.51	49.41	41.23
	(4.65)	(2.50)	(5.30)	(2.00)
Kampot	45.63	46.36	38.72	37.86
	(6.09)	(3.40)	(6.27)	(3.47)
Kandal	46.25	47.88	48.16	49.04
	(5.38)	(2.01)	(4.65)	(2.07)
Koh Kong	55.03	45.30	42.72	44.93
	(4.81)	(3.51)	(4.83)	(3.05)
Kratie	55.22	52.72	49.75	52.21
	(6.29)	(1.61)	(4.68)	(2.14)
Mondolkiri	43.89	42.63	48.44	36.96
	(6.68)	(2.47)	(7.06)	(2.72)
Phnom Penh	25.59	32.76	35.04	34.23
	(4.75)	(1.66)	(4.95)	(1.87)
Preah Vihear	46.71	49.33	35.74	46.61
	(8.37)	(2.17)	(8.91)	(2.17)
Prey Veng	51.22	51.98	56.82	44.74
	(5.82)	(1.91)	(5.96)	(1.92)
Pursat	46.28	46.19	46.34	45.72
	(4.40)	(2.19)	(4.99)	(2.60)
Rotanakiri	59.99	51.45	56.60	46.35
	(4.46)	(2.29)	(4.98)	(2.61)
Siem Reap	52.60	44.36	51.77	43.46
	(4.85)	(2.79)	(4.61)	(2.00)
Krong Preah	39.84	43.73	43.81	41.43
Sihanouk	(11.96)	(1.67)	(11.36)	(2.02)
Stung Traeng	46.79	50.39	58.54	46.27
	(12.02)	(1.84)	(11.23)	(2.18)
Svay Rieng	51.33	51.44	45.86	45.21
	(5.66)	(2.36)	(5.72)	(1.93)
Takeo	42.09	47.98	39.92	44.92
	(4.84)	(2.01)	(4.85)	(2.41)

(Continued)

TABLE 5. Continued

	$P^{0,(1)}$ (% stunted)		$P^{0,(2)}$ (% underweight)	
Province	Survey only Estimate (standard error)	Small-area estimation Estimate (standard error)	Survey only Estimate (standard error)	Small-area estimation Estimate (standard error)
Otdar Mean Chey	36.07	46.22	35.29	42.31
	(8.82)	(2.86)	(11.22)	(2.63)
Krong Keb	36.36	30.18	27.27	37.95
	(10.64)	(2.11)	(10.77)	(2.75)
Krong Pailin	36.36	23.10	27.27	29.39
	(11.13)	(2.41)	(10.27)	(2.28)

Note: The standard errors for the survey-only estimates were calculated by a 100-time two-stage bootstrapping simulation.

Source: Author's analysis based on the Cambodia Demographic and Health Survey for 2000 and the Cambodia National Population Census for 1998.

aggregated unit is used as a cluster, so the direction of change for the estimated standard errors is ambiguous.

This does not mean that the clustering unit is unimportant. Among the various clustering units considered, the standard errors for commune-level estimates are the highest for all quartiles when the clustering unit is the commune. Similarly, the standard errors for district-level estimates are highest when the clustering unit is the district. Hence, the standard errors based on village-level clustering may be underestimated by about 15 percent if the clustering occurs at a different level. Even though the most appropriate results cannot be determined by merely looking at table 4, conservative estimates of the standard errors can be obtained by inflating the standard errors for small-area estimates with commu-level clustering by 15 percent. Even with this correction, the estimates remain reasonably accurate at the commune level.

Fifth, the standard errors for the small-area estimation are smaller than those for the survey-only estimates. This is because the small-area estimation does not suffer from sampling error, since the census covers all the individuals or households of interest in the country. However, the small-area estimates suffer from both idiosyncratic and model errors, which the survey-only estimates are not subject to. As a result of these different sources of error, whether the small-area estimates have smaller standard errors than the survey-only estimates depends on the relative magnitudes of these sources. In this application small-area estimation brings about a large gain in accuracy in the Coastal ecozone because its survey sample size is relatively small.

Small-area estimates are often found to have smaller standard errors than survey-only estimates in poverty mapping as well. For example, Elbers and others (2003) report survey-only estimates and small-area estimates of the headcount index of poverty for 31 regions in three countries, and the standard error is smaller for the small-area estimates in 27 of them. Thus, like the current study, their study shows that the small-area estimation can often improve aggregate estimates.

In addition to the comparisons at the ecozone and province levels, the small-area estimates are also compared with the survey-only estimates for more than 400 survey villages. This is a meaningful comparison because the survey contains a representative sample of households for these villages. To this end, standard errors are calculated for the survey-only estimates by bootstrapping, to take into account that children in the same household are included in the sample together. Each village averages only eight children, so the survey-only estimates are very noisy, with an average standard error of over 25 percent for both stunting and underweight. The small fraction of villages in which the prevalence of undernourished children is 0 or 1 was excluded from this calculation because reasonable standard errors cannot be obtained without making some arbitrary assumptions.

Because idiosyncratic error is a major source of error for the small-area estimation at a disaggregated level, the correlation between the estimation errors for the survey-only and small-area estimates are unlikely to be high. Thus, assuming that the two estimation errors are uncorrelated, a Z-statistic can be computed for each village by taking their difference and dividing it by the square root of the sum of their variances. Under this assumption the proportions of the villages with a Z-statistic for stunting of over 1 and over 2 in absolute value are about 3 percent and 17 percent. These proportions are less than the 5 percent and 32 percent expected from the standard normal distribution. The results for underweight are qualitatively similar.

As a robustness check, the Z-statistics under the alternative assumption of a perfectly positive correlation were computed. Even under this conservative assumption, the proportion of villages with a Z-statistic for stunting of over 1 and over 2 in absolute value are about 8 percent and 32 percent, which differ little from what is expected from the standard normal distribution.

These results may merely indicate that the survey-only estimates are too noisy. Thus, the correlation between the survey-only and the small-area estimates for survey villages are also examined. Because the estimates are noisy at the village level, the correlations are not strong (0.29 for stunting and 0.18 for underweight), but they are significant. All these results indicate that the small-area estimates for the survey villages are consistent with the corresponding survey-only estimates.

So far, the discussion has been based on the cutoff Z-score of -2. Although this cutoff is standard, changing it allows more or less severe undernutrition to be examined. Therefore, the robustness of the methodology with respect to the choice of the cutoff Z-score was checked by setting the threshold at -1, -1.5, -2 (baseline), -2.5, and -3. The differences between the survey-only and the small-area estimates are generally small. For example, for the Coastal ecozone the small-area estimates of the prevalence ($P^{0,(1)}$) and gap ($P^{1,(1)}$) of stunting are within two standard errors of the survey-only estimates for all five thresholds.

The importance of including the intrapersonal correlation is also evaluated. To this end, a small-area estimation assuming no intrapersonal correlation was conducted and the results compared with the baseline small-area estimates. Under the assumption of no intrapersonal correlation, the generalized least squares estimation was run with the intrapersonal correlation equal to 0, so that the regressions are effectively run separately. Further, each component of the disturbance terms was drawn independently across the indicators. Hence, the small-area estimation under the assumption of no intrapersonal correlation is equivalent to carrying out the small-area estimation separately for each anthropometric indicator.

The aggregate estimates and their standard errors under the assumption of no intrapersonal correlation are very similar to those reported in table 4. However, the intrapersonal correlation affects the estimated correlation between the prevalence of stunting and of underweight at aggregate levels. The survey-only estimate of the correlation between stunting and underweight is compared with the corresponding small-area estimate at the district and province levels but not at the commune level because about 70 percent of communes have no observations at all in the CDHS data. In contrast, the CDHS contains some observations for more than 90 percent of districts and all provinces in Cambodia. And while the survey-only estimates of the prevalence of undernutrition are unreliable because of the small sample size, sampling does not systematically bias the survey-only estimate of correlation.

The survey-only estimate of the correlation between the district-level prevalence of stunting and underweight is 62.6 percent with bootstrap standard errors of 6.9 percent, and the survey-only estimate of the correlation between the province-level prevalence of stunting and underweight is 69.6 percent with bootstrap standard errors of 12.3 percent. The corresponding figures for the baseline small-area estimates at the district level are 53.7 percent with a standard error of 6.6 percent and at the province level are 84.3 percent with a standard errors of 6.8 percent. Therefore, the baseline estimates of district- and province-level correlation are not significantly different from the survey-only estimate. But when the intrapersonal correlation is ignored in the small-area estimation, the corresponding figures are 26.7 percent with a standard error of 6.2 percent at the district level and 27.7 percent with a standard error of 15.0 percent at the province level.

This indicates that the correlation between two aggregate measures can be severely underestimated if they are calculated from two sets of small-area estimates produced independently. This occurs because small-area estimates of two indicators are subject to random errors, which may be correlated across indicators. Such correlation is not appropriately captured when the two sets of small-area estimates are produced independently. Therefore, it cannot be concluded that two aggregate indicators have no correlation solely because the small-area estimates of these aggregate indicators have no correlation.

This is an important point because the correlations at an aggregate level may have important policy implications. In the example here, if the intrapersonal correlation is not taken into account, the prevalence of stunting and underweight would appear to have only a weak correlation at an aggregate level. Thus, it would appear that two completely different sets of measures may be necessary to address medium- and long-term undernutrition. Because such a conclusion could lead to inappropriate policies, it is important to address the intrapersonal correlations if the correlations across aggregate indicators are of interest.

While the small-area estimates can be provided in a table, it is also useful to convert them into maps using a geographic information system (maps 1–5). Although there is no simple geographic pattern of undernutrition, the areas

MAP 1. Commune-Level Prevalence of Stunting for Children under Age 5 in Cambodia

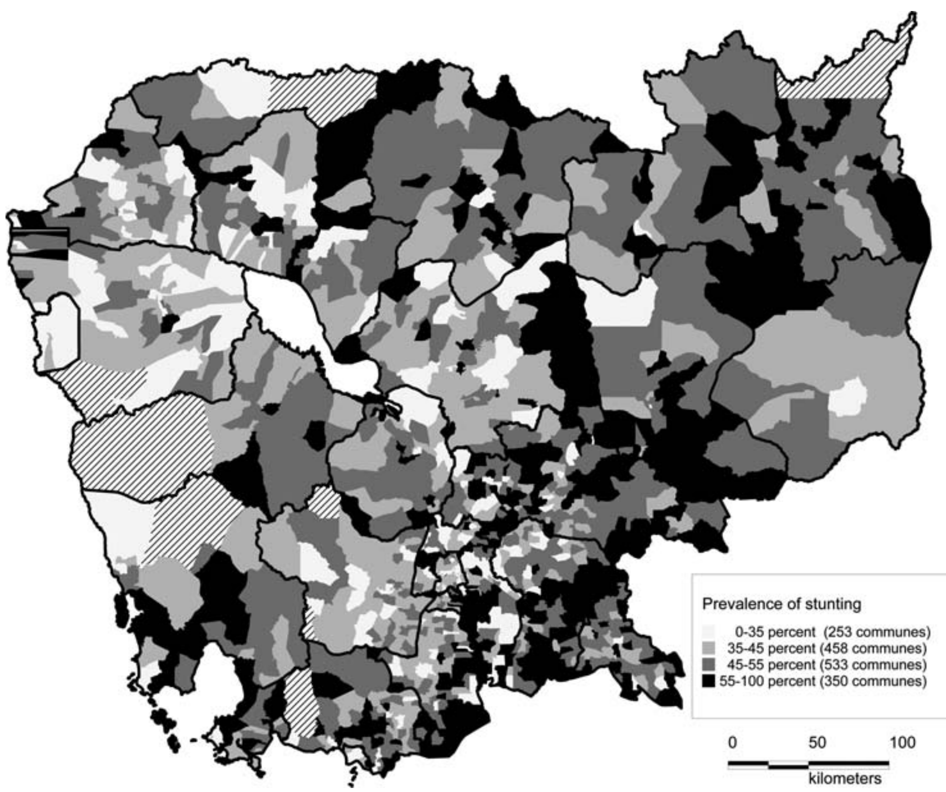

Source: Author's analysis based on the Cambodia Demographic and Health Survey for 2000 and the Cambodia National Population Census for 1998.

MAP 2. Prevalence of Stunting and Underweight at the Commune Level in Cambodia

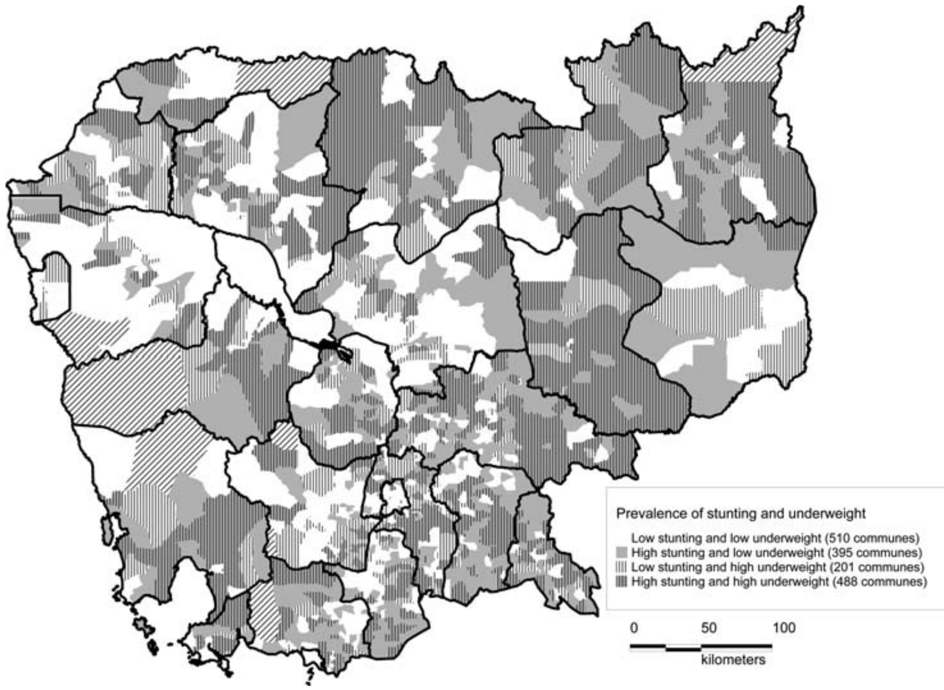

Note: High and low indicate above and below 45 percent.

Source: Author's analysis based on the Cambodia Demographic and Health Survey for 2000 and the Cambodia National Population Census for 1998.

with stunting and underweight generally exhibit similar patterns, with a commune-level correlation of 0.37.

The prevalence of both stunting and underweight is high in the most densely populated parts of Cambodia surrounding Phnom Penh, the provinces of Kandal, Prey Veng, Svay Rieng, Kampong Cham, and Kampong Chhnang (see map 3). Although these provinces are geographically close to each other, the causes of undernutrition may be different. For example, the prevalence of diarrhea for children under age 5 is 29.7 percent in Kandal but only 3.1 percent in Prey Veng. Similarly, the prevalence of fever, a primary manifestation of malaria and other acute infections in children, is 46.8 percent in Kandal and 4.0 percent in Prey Veng (Cambodian National Institute of Statistics, Cambodian Directorate General for Health, and ORC Macro 2001). However, the poverty rates in Kandal and Prey Veng are estimated at 18.4 percent and 53.1 percent (Fujii 2006). This suggests that undernutrition in Prey Veng is likely to be driven mainly by poverty, or more specifically lack of caloric intake, while infectious diseases seem to be important causes of undernutrition in Kandal.

MAP 3. Prevalence of Stunting and Underweight around Phnom Penh

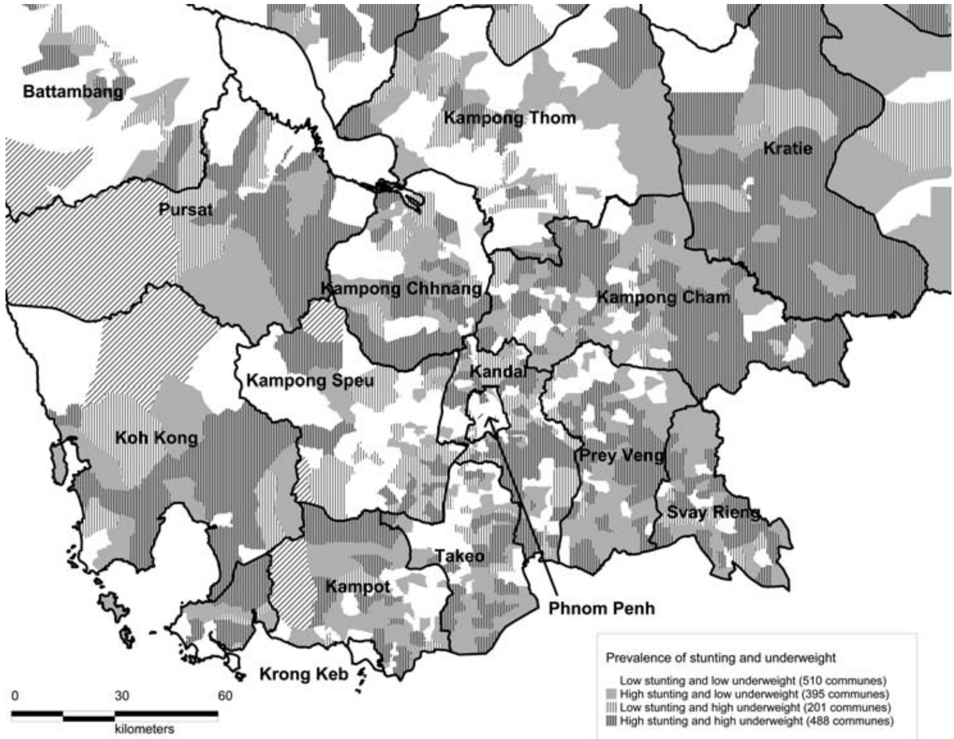

Note: High and low indicate above and below 45 percent.
Source: Author's analysis based on the Cambodia Demographic and Health Survey for 2000 and the Cambodia National Population Census for 1998.

Map 3 also illustrates some noticeable differences between stunting and underweight. For example, most areas of the Kampot province have a high prevalence of stunting, but only a few areas have a high prevalence of underweight. This means that nutrition status has improved over time, which may be due partly to improvement in road access. Low levels of stunting and high levels of underweight would reflect a recent aggravation of nutrition status, possible causes of which include increased incidence of malaria and diarrhea and acute food shortage due to natural disaster.

Maps 1, 2, and 3, while presented in a user-friendly format, ignore the fact that the commune-level estimates are subject to statistical errors. Map 4 provides an alternative representation of the small-area estimates in which the map is colored according to the difference between the commune-level estimate and the national average divided by the standard error of the commune-level estimate. This representation allows the communes that are significantly more or less undernourished than the national average to be idenfitied.

More than half the communes are significantly more or less undernourished than the national average. This is clearly a substantial improvement from the

MAP 4. Difference in the Prevalence of Stunting between the National Average and the Commune-Level Estimate as Measured by the Number of Small-Area Estimation Standard Errors

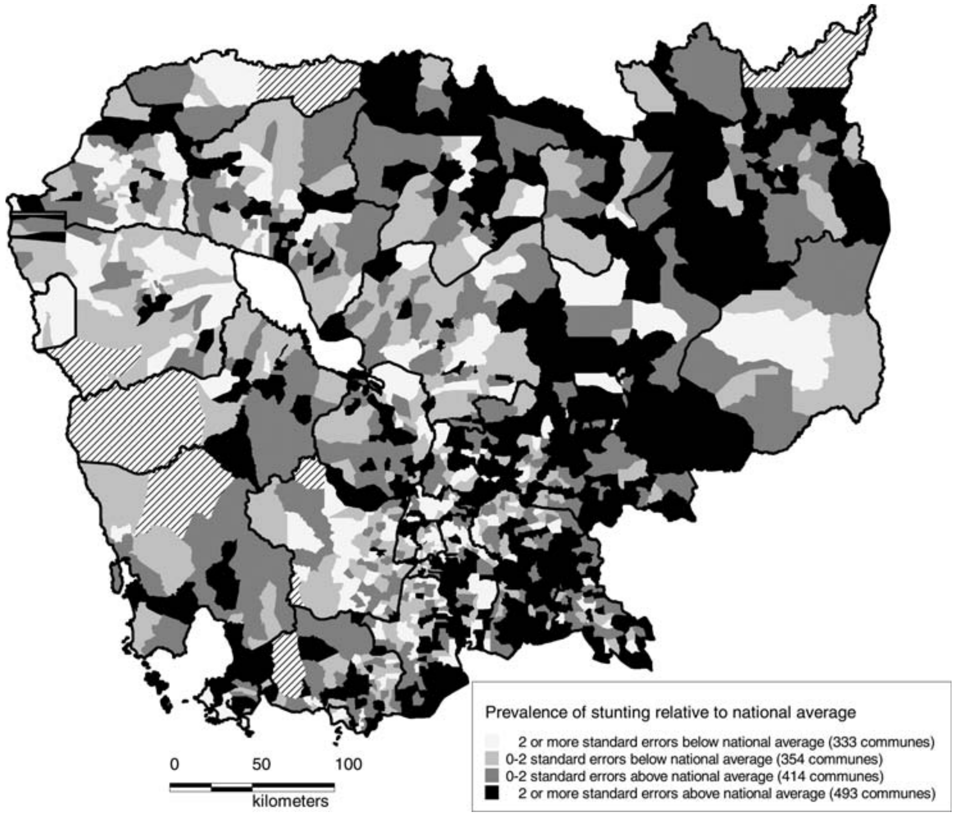

Source: Author's analysis based on the Cambodia Demographic and Health Survey for 2000 and the Cambodia National Population Census for 1998.

survey-only estimates. It could be argued that the reference point should be the ecozone-level estimates instead of the national average because the survey-only estimates are reasonably accurate at the ecozone level. Even when the ecozone level is used as the reference point, more than 45 percent of the communes have an estimate of the prevalence of stunting that is significantly different from the ecozone average.

Another useful way to present the results is to depict the density of undernourished children. As shown in map 5, the picture is completely different from map 1 because of the high population density in Phnom Penh and the surrounding provinces as well as the areas around the Tonle Sap Lake.

Different representations can give different impressions about the situation of undernutrition. No representation is better than others under all circumstances because different representations describe different aspects of undernutrition.

MAP 5. The Density of Stunted Children in Cambodia

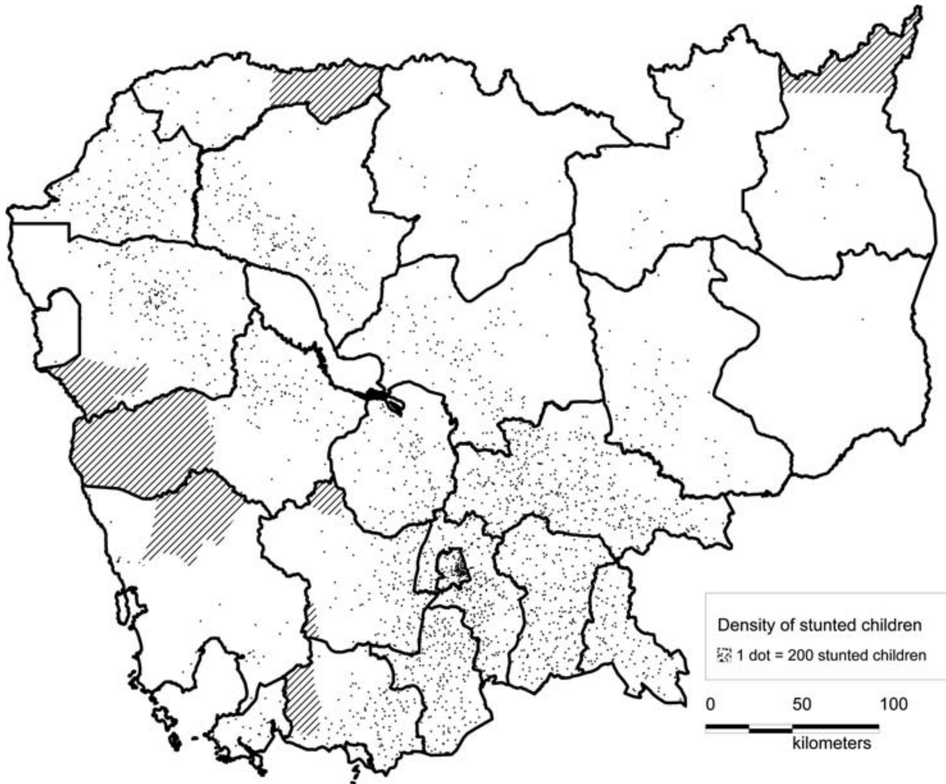

Source: Author's analysis based on the Cambodia Demographic and Health Survey for 2000 and the Cambodia National Population Census for 1998.

It is thus important to choose the representation that suits the purpose and the audience of the map. The density representation is more appropriate for policymakers concerned about the absolute number of undernourished children. Consider the construction of health clinics. If undernourished children in the neighborhood of a health clinic would benefit from the clinic, the number of undernourished children would be the appropriate yardstick, not the prevalence of undernutrition.

Cambodian health clinics are likely to help improve the nutrition status of children in areas such as Kandal, where infectious disease is a likely cause of undernutrition and where the distance to the nearest health facility is far. In fact, 60.1 percent of women in Kandal reported that distance to the nearest health facility is a big problem for personal access to health care. This contrasts with 12.9 percent in Prey Veng, where undernutrition is probably due to poverty.

Policymakers may also be concerned about the prevalence of undernutrition in each location. If food or micronutrient supplements were distributed at the commune level to improve the nutrition status of children, policymakers should examine the prevalence of undernutrition to minimize the leakage of resources to well nourished children (see maps 1 and 4). Point estimates would be more useful when a large number of communes are targeted because the standard error for each commune is not particularly important. But the deviation from the national average over standard error would be an appropriate representation when policymakers want to select only a few communes for a nutrition program.

V. Conclusion

Estimates of the prevalence of child undernutrition were previously available only at the CDHS stratum level. Stratum-level estimates often mask great disparities in the prevalence of undernutrition within the stratum. Unless there are strata with an extremely high prevalence of undernutrition, targeting based on stratum-level estimates is unlikely to capture many of the undernourished children, and misallocation of resources is likely.

To overcome the problem of limited data, a methodology was developed to estimate the prevalence of child undernutrition at the level of small geographic areas. The estimates of the prevalence of child undernutrition in Cambodia were also disaggregated from the 17 CDHS strata into 1,594 communes. The Elbers, Lanjouw, and Lanjouw (2002, 2003) small-area estimation technique was extended to jointly estimate multiple indicators and allow for a richer structure of error terms. This is a crucial step to address issues unique to nutrition indicators. Although this methodology was applied to the Cambodian data, it can be easily applied to other countries where census data and survey data with an anthropometric component are available.

At the ecozone and province levels the small-area estimates closely match the survey estimates and generally have lower standard errors than the survey estimates. Further, the magnitude of the standard errors of the commune-level estimates of the prevalence of undernutrition is largely comparable to that of the survey estimates at the stratum level. The estimation results are robust to the changes in the level of clustering and yield good estimates of the correlation of aggregate indicators after accounting for the intrapersonal correlations.

The small-area estimates can be easily provided in a table, but presenting them in maps allows policymakers to see areas of severe undernutrition and to formulate targeting policies. Several alternative representations of the small-area estimates, each with a specific purpose, were discussed.

The nutrition maps overlaid with other maps help identify possible causes of undernutrition in different locations. This in turn provides policymakers with valuable information on the appropriate design of child nutrition programs. For example, consider a nutrition map combined with a map of areas affected by natural disasters, such as flood and drought. Natural disasters affect agricultural output severely, and most Cambodian farmers have limited means to cope with disasters. As a result, the nutrition status of children in disaster-affected areas can be negatively affected. The overlaid map helps policymakers identify such areas. Because food shortage is the cause of undernutrition in these areas, food relief there may be appropriate.[7]

Similarly, nutrition maps overlaid with the map of the prevalence of diseases such as malaria and diarrhea can help identify areas where the primary cause of undernutrition is disease. One can further superimpose the map of the location of the health clinics to identify where the prevalence of undernutrition and disease is high and no health clinic is available nearby. In such areas building local health clinics may be more efficient at reducing child undernutrition than delivering food relief, for example. Thus, overlaid maps are helpful for policymakers not only in choosing the target areas but also in determining the appropriate interventions.

Appendix. Technical Details of the Methodology

This appendix presents the technical details of small-area estimation. To this end, some additional notations are introduced. The set of all clusters is denoted by $\mathcal{C}$, the set of all households in cluster $c(\in \mathcal{C})$ is denoted by $\mathcal{H}_c$, and the set of all children under age 5 in household $h(\in \mathcal{H}_c)$ is denoted by $\mathcal{I}_{ch}$. Each household and each child belong to exactly one cluster and one household respectively. $\mathrm{C} \equiv \#(\mathcal{C})$, $\mathrm{H}_c \equiv \#(\mathcal{H}_c)$, and $\mathrm{I}_{ch} \equiv \#(\mathcal{I}_{ch})$, where $\#(\cdot)$ is the counting measure. Each cluster has weight w_c , which is normalized so that $\sum_c w_c = 1$.

The variances of $\eta_c^{(k)}$, $\epsilon_{ch}^{(k)}$, and $\delta_{chi}^{(k)}$ are denoted by $(\sigma_\eta^{(k)})^2$, $(\sigma_{\epsilon,\ ch}^{(k)})^2$, and $(\sigma_\eta^{(k)})^2$. The subscript ch is needed for the household-specific effect because of its heteroskedasticity. The correlation of individual effects across indicators (intrapersonal correlation) is denoted by $\rho^{(k,l)} = \frac{\sigma_\delta^{(k,l)}}{\sigma_\delta^{(k)}\sigma_\delta^{(l)}}$, where $\sigma_\delta^{(k,\mathrm{l})}$ is the covariance of the individual effects between the k-th and l-th anthropometric indicators.

To run a generalized least squares regression, the variance-covariance matrix Ω of the disturbance term u must be estimated. To this end, an ordinary least squares regression of equation 1 is run for each k. The ordinary least squares regression residuals $\hat{u}_{chi}^{(k)}$ are then used to estimate various components of Ω.

To estimate $(\sigma_\delta^{(k)})^2$ and $\sigma_\delta^{(k,l)}$, the following equations, which are straightforward to verify, are used:

$$\text{(A1)} \qquad (\sigma_\delta^{(k)})^2 = E\left[\sum_{c\in\tilde{C}} \frac{\tilde{w}_c}{\tilde{\mathrm{H}}_c} \sum_{h\in\tilde{\mathcal{H}}_c} \sum_{i\in\mathcal{I}_{ch}} \frac{(u_{chi}^{(k)} - u_{ch\cdot}^{(k)})^2}{\mathrm{I}_{ch} - 1}\right]$$

$$\text{(A2)} \qquad \sigma_\delta^{(k,l)} = E\left[\sum_{c\in\tilde{C}} \frac{\tilde{w}_c}{\tilde{\mathrm{H}}_c} \sum_{h\in\tilde{\mathcal{H}}_c} \sum_{i\in\mathcal{I}_{ch}} \frac{(u_{chi}^{(k)} - u_{ch\cdot}^{(k)}) \cdot (u_{chi}^{(l)} - u_{ch\cdot}^{(l)})}{I_{ch} - 1}\right],$$

7. In fact, the World Food Programme in Cambodia has created and used a preliminary version of nutrition maps overlaid with a map of areas affected by flood and drought (Fujii 2007b).

where $\tilde{\mathcal{H}}_c \equiv \{h \in \mathcal{H}_c | I_{ch} > 1\}$, $\tilde{\mathrm{H}}_c \equiv \#\{\tilde{\mathcal{H}}_c\}$, $\tilde{\mathcal{C}} \equiv \{c \in \mathcal{C} | \tilde{\mathrm{H}}_c > 0\}$, and $\tilde{w}_c \equiv \frac{w_c}{\sum_{c' \in \tilde{\mathcal{C}}} w_{c'}}$. A tilde is used to make it clear that a subset of the sample is used for some parts of the estimation.

Consistent estimators $(\hat{\sigma}_\delta^{(k)})^2$ and $\hat{\sigma}_\delta^{(k,l)}$ are obtained by removing the expectations operator on the right side and by replacing u with $\hat{u}$ in the respective equations. The intrapersonal correlation is simply estimated at $\hat{\rho}^{(k,l)} = \frac{\hat{\sigma}_\delta^{(k,l)}}{\hat{\sigma}_\delta^{(k)} \hat{\sigma}_\delta^{(l)}}$.

To estimate the variance of the cluster-specific effect, the following equation is used:

$$\text{(A3)} \qquad (\sigma_\eta^{(k)})^2 = E\left[\frac{\sum_{c \in C} w_c \mathrm{H}_c (u_{c\cdot\cdot}^{(k)})^2 - \sum_{c \in C} \frac{w_c}{\mathrm{H}_c} \sum_{h \in \mathcal{H}_c} (u_{ch\cdot}^{(k)})^2}{\sum_{c \in C} w_c (\mathrm{H}_c - 1)}\right].$$

Proof of equation (A3):

$$E[(u_{ch\cdot}^{(k)})^2] = (\sigma_\eta^{(k)})^2 + (\sigma_{\epsilon,ch}^{(k)})^2 + \frac{1}{I_{ch}} (\sigma_\delta^{(k)})^2,$$

$$E[(u_{c\cdot\cdot}^{(k)})^2] = (\sigma_\eta^{(k)})^2 + \frac{1}{\mathrm{H}_c^2} \sum_{h \in \mathcal{H}_c} (\sigma_{\epsilon,ch}^{(k)})^2 + \frac{1}{\mathrm{H}_c^2} \left(\sum_{h \in \mathcal{H}_c} \frac{1}{I_{ch}}\right) (\sigma_\delta^{(k)})^2.$$

Hence,

$$\begin{aligned}
\sum_{c \in \mathcal{C}} \frac{w_c}{\mathrm{H}_c} \sum_{h \in \mathcal{H}_c} E\left[(u_{ch\cdot})^2\right] &= \sigma_\eta^2 + \sum_{c \in \mathcal{C}} w_c \mathrm{H}_c \left(\frac{1}{\mathrm{H}_c^2} \left(\sum_{h \in \mathcal{H}_c} \sigma_{\epsilon,ch}^2\right) + \frac{1}{\mathrm{H}_c^2} \left(\sum_{h \in \mathcal{H}_c} \frac{1}{I_{ch}}\right) \sigma_\delta^2\right) \\
&= \sigma_\eta^2 + \sum_{c \in \mathcal{C}} w_c \mathrm{H}_c \left(E\left[(u_{c\cdot\cdot})^2\right] - \sigma_\eta^2\right) \\
&= (1 - \sum_{c \in \mathcal{C}} w_c \mathrm{H}_c) \sigma_\eta^2 + \sum_{c \in C} w_c \mathrm{H}_c E\left[(u_{c\cdot\cdot})^2\right].
\end{aligned}$$

Solving for σ_η^2 yields equaton (A3).

As with $(\hat{\sigma}_\delta^{(k)})^2$ and $\hat{\sigma}_\delta^{(k,l)}$, $(\sigma_\eta^{(k)})^2$ is estimated by removing the expectation operator and replacing u with $\hat{u}$ in equation (A3). The estimate is censored at 0 to guarantee the non-negativity of the variance.

To estimate the variance of the household-specific effect, the sum $(s_{ch}^{(k)})^2$ $(\equiv (\sigma_{\epsilon,\ ch}^{(k)})^2 + (\sigma_\delta^{(k)})^2)$ is used first. The following formula is used for $c \in \tilde{C}^* (\equiv \{\psi \in \mathcal{C} | \mathrm{H}_\psi > \in\})$:

$$\text{(A4)} \quad E\left[\frac{\mathrm{H}_c \cdot (u_{ch\cdot}^{(k)} - u_{c\cdot\cdot}^{(k)})^2}{\mathrm{H}_c - 2} - \frac{\sum_{h' \in \mathcal{H}_c} (u_{ch'\cdot}^{(k)} - u_{c\cdot\cdot}^{(k)})^2}{(\mathrm{H}_c - 1)(\mathrm{H}_c - 2)}\right] + \frac{I_{ch} - 1}{I_{ch}} (\sigma_\delta^{(k)})^2 = (s_{ch}^{(k)})^2.$$

Proof of equation (A4):

$$E[(u_{ch\cdot} - u_{c\cdot\cdot})^2] = \frac{\mathrm{H}_c - 2}{\mathrm{H}_c}\left\{\frac{\sigma_\delta^2}{I_{ch}} + \sigma^2_{\epsilon,ch}\right\} + \frac{1}{\mathrm{H}_c^2}\left\{\sum_{h'\in\mathcal{H}_c}\left(\frac{\sigma_\delta}{I_{ch'}} + \sigma^2_{\epsilon,ch'}\right)\right\},$$

$$E[\sum_{h'\in\mathcal{H}_c}(u_{ch'\cdot} - u_{c\cdot\cdot})^2] = \frac{\mathrm{H}_c - 1}{\mathrm{H}_c}\left\{\sum_{h'\in\mathcal{H}_c}\left(\frac{\sigma_\delta^2}{\mathcal{I}_{ch'}} + \sigma^2_{\epsilon,ch'}\right)\right\}.$$

It is straightforward to show the first equation above. The second equation can be derived by summing the first equation over the households in each cluster. From these two equations, the following holds for all the households in $\tilde{\mathcal{C}}*$:

$$\sigma^2_{\epsilon,ch} = \frac{\mathrm{H}_c}{\mathrm{H}_c - 2}E[(u_{ch\cdot} - u_{c\cdot\cdot})^2] - \frac{1}{(\mathrm{H}_c - 1)(\mathrm{H}_c - 2)}E[\sum_{h'\in\mathcal{H}_c}(u_{ch'\cdot} - u_{c\cdot\cdot})^2] - \frac{\sigma_\delta^2}{I_{ch}}.$$

Adding σ_δ^2 to both sides of the equality and arranging the terms yields equation (A4).

Let $(\hat{s}^{(k)}_{ch})^2$ be the left side of equation (A4) above with the expectation operator removed and with u and σ_δ replaced by $\hat{u}$ and $\hat{\sigma}_\delta$, respectively.[8] For the heteroskedastic model the following logistic heteroskedastic model, similar to the one in Elbers, Lanjouw, and Lanjouw (2002):

$$\text{(A5)} \qquad \ln\frac{(\hat{s}^{(k)}_{ch})^2 - B^{(k)}}{A^{(k)} + B^{(k)} - (\hat{s}^{(k)}_{ch})^2} = [\mathbf{z}^{(k)}_{ch}]^T\alpha^{(k)} + \tau^{(k)}_{ch},$$

where $A^{(k)}$ and $B(k)$ are the maximum and minimum of $(s^{(k)}_{ch})^2$, $\mathrm{z}^{(k)}_{ch}$ is a vector of household characteristics, $\alpha^{(k)}$ is a heteroskedastic regression coefficient, and $\tau^{(k)}_{ch}$ is the residual term. The feature of this formulation is that $(\hat{s}^{(k)})^2$ is both upper- and lower-bounded. Using $B^{(k)}_*= \min\{0, 1.05 \cdot \min_{ch}\{(\hat{s}^{(k)}_{ch})^2\}\}$ and $A^{(k)}_*= 1.05 \cdot (\max_{ch}\{(\hat{s}^{(k)}_{ch})^2\}-B^{(k)}_*)$, an ordinary least squares regression is run to obtain $\hat{\alpha}^{(k)}$. In the application to Cambodia, $B^{(k)}_* = 0$ in all ecozones, so the heteroskedastic model is identical to the one used in Elbers, Lanjouw, and Lanjouw (2002). However, in general, an explicit lower bound is needed because $(\hat{s}^{(k)}_{ch})^2$ is not guaranteed to be positive for all c and h.

Using the delta method yields the following estimate of $(\sigma^{(k)}_{\epsilon,\,ch})^2$:

$$(\hat{\sigma}^{(k)}_{\epsilon,ch})^2 = \max\left\{0, \left[\frac{D^{(k)}}{1 + D^{(k)}} + \frac{D^{(k)}(1 - D^{(k)})(\hat{\sigma}^{(k)}_\tau)^2}{2(1 + D^{(k)})^3}\right]A^{(k)}_* + B^{(k)}_* - (\hat{\sigma}^{(k)}_\delta)^2\right\}.$$

(A6)

where $D^{(k)} \equiv \exp([\mathbf{z}^{(k)}_{ch}]^T\,\hat{\alpha}^{(k)})$, and $(\hat{\sigma}^{(k)}_\tau)^2$ is the estimated variance of τ_{ch}. The

8. When $(\hat{\sigma}^{(k)}_\eta)^2 = 0, E[(u^{(k)}_{ch\cdot})^2] = (\sigma^{(k)}_{\epsilon,ch}) + \frac{(\sigma^{(k)}_\delta)^2}{I_{ch}}$ is used instead.

max{·, ·} function is introduced to ensure the non-negativity of $(\hat{\sigma}^{(k)}_{\varepsilon,ch})^2$. The consequence of this is that $(\sigma^{(k)}_{\varepsilon,ch})^2$ may be upward biased. Hence, the standard errors for the estimates of nutrition measures at the level of small geographic areas are conservative. The variance-covariance matrix Ω can be estimated and a (feasible) generalized least squares regression performed to obtain the regression coefficients $\hat{\beta}$ and $\widehat{Var}[\hat{\beta}]$ for all indicators simultaneously. The distributions of η_c, ϵ_{ch}, and δ_{chi} are approximated by the empirical distributions of $u_{c\,.\,.}$, $(u_{ch\,.} - u_{c\,.\,.})$, and $(u_{chi} - u_{ch\,.})$, respectively, standardized to have mean 0 and a unit standard error.

With these estimates, the simulation can proceed. Let R be the number of simulations, which must be sufficiently large to make the computational errors sufficiently small. In this study $R = 100$. In the r-th simulation where $r \in \{1, 2, \ldots, R\}$, the following parameters are needed: $\tilde{\beta}^{(k)}_{(r)}$, $\tilde{\alpha}^{(k)}_{(r)}$, $(\tilde{\sigma}_{\tau,(r)})^2$, $(\tilde{\sigma}^{(k)}_{\delta,(r)})^2$, $(\tilde{\sigma}^{(k,l)}_{\delta,(r)})^2$, $(\tilde{\sigma}^{(k)}_{\eta,(r)})^2$, $A^{(k)}_{*,(r)}$, $B^{(k)}_{*,(r)}$ for $\forall k \neq$ l. First, $\tilde{\alpha}_{(r)}$ and $\tilde{\beta}_{(r)}$ are randomly drawn from the normal distribution with mean $\hat{\alpha}$ and $\hat{\beta}$, and variance-covariance matrix $\widehat{Var}[\hat{\alpha}]$ and $\widehat{Var}[\hat{\beta}]$ respectively. For the rest of the parameters, a two-stage bootstrap sample of $\hat{u}$ is created in each round of simulation, and the parameters are computed using the bootstrapping sample.[9] It is straightforward to calculate $(\tilde{\sigma}^{(k)}_{\varepsilon,ch,(r)})^2$ from equation (A6).

For each cluster, household, and individual in the census the standardized cluster-, household-, and individual-specific effects are drawn. Each component of the disturbance term is drawn jointly for multiple indicators to capture the correlation of the disturbance term across indicators. Letting the standardized components of the disturbance terms drawn in the r-th simulation be $\tilde{\eta}^{(k)}_{c,(r)}$, $\tilde{\varepsilon}^{(k)}_{ch,(r)}$, and $\tilde{\delta}^{(k)}_{chi}$ respectively, the k-th imputed anthropometric indicator for the individual in the census in the r-th simulation is:[10]

$$\tilde{y}^{(k)}_{chi,(r)} = \mathbf{x}^{(k)}_{chi}\tilde{\beta}^{(k)}_{(r)} + \tilde{\eta}^{(k)}_{c,(r)} \cdot \tilde{\sigma}_{\eta,(r)} + \tilde{\epsilon}^{(k)}_{ch,(r)} \cdot \tilde{\sigma}_{\epsilon,ch,(r)} + \tilde{\delta}^{(k)}_{chi} \cdot \tilde{\sigma}_{\delta,(r)} \tag{A7}$$

Suppose that an estimate of an aggregate index $\mathbf{W}^{(k)}_{\nu} \equiv \mathbf{W}(\{y^{(k)}_i\}_{i\in\nu})$ is needed, such as the prevalence of undernutrition for a set of individuals ν. An estimate $\tilde{\mathbf{W}}^{(k)}_{(r),\nu} = \mathbf{W}(\{\tilde{y}^{(k)}_{i,(r)}\}_{i\in\nu})$ can be obtained in the r-th simulation using $\tilde{y}^{(k)}_{chi,(r)}$. This is subject to the particular realization of the model error and idiosyncratic error in the r-th simulation. However, robust point estimates and standard errors can be derived by taking the mean and the standard deviation of $\tilde{\mathbf{W}}^{(k)}_{(r),V}$ over r. For example, the point estimate of the prevalence of stunting for

9. Another possible implementation is to draw $\tilde{\alpha}$ and $\tilde{\beta}$ from the bootstrapping sample.

10. To eliminate extreme values, census observations for which the point estimate $\mathbf{x}^{(k)}_{chi}\hat{\beta}^{(k)}$ is not in the range of $y^{(k)}$ for at least one indicator are dropped. In each simulation $\tilde{y}^{(k)}_{chi,(r)}$ was censored at the minimum and maximum observed in the survey. $D^{(k)}$ was censored at the minimum and maximum.

commune κ is:

$$\hat{\mathrm{P}}_{\kappa}^{0,(1)} = \frac{1}{R \cdot \sum_{c \in \mathcal{C}_\kappa} \sum_{h \in \mathcal{H}_c} I_{ch}} \cdot \sum_{r=1}^{R} \sum_{c \in \mathcal{C}_\kappa} \sum_{h \in \mathcal{H}_c} \sum_{i \in \mathcal{I}_{ch}} \mathrm{Ind}(\tilde{y}_{chi,(r)}^{(1)} < z^{(1)}).$$

References

Alderman, H. 2000. "Anthropometry." In M. Grosh, and P. Glewwe, eds, *Designing Household Survey Questionnaires for Developing Countries: Lessons from Ten Years of LSMS Experience Volume One*. Oxford, UK: Oxford University Press.

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.

Cambodian National Institute of Statistics, Cambodian Directorate General for Health, and ORC Macro. 2001. *Cambodia Demographic and Health Survey 2000*. Phnom Penh and Calverton, Md.: Cambodian National Institute of Statistics, Cambodian Directorate General for Health, and ORC Macro.

Curtis, S.L., and M. Hossain 1998. "The Effects of Aridity Zone on Child Nutritional Status." West Africa Spatial Analysis Prototype Explanatory Analysis. Macro International, Calverton, Md.

de Onis, M., E.A. Frongillo, and M. Blössner. 2000. "Is Malnutrition Declining? An Analysis of Changes in Levels of Child Malnutrition since 1980." *Bulletin of the World Health Organization* 78: 1222–33.

Demombynes, G., and B. Özler 2005. "Crime and Local Inequality in South Africa." *Journal of Development Economics* 76(2): 265–92.

Demombynes, G., C. Elbers, J.O. Lanjouw, and P. Lanjouw 2007. "How Good a Map? Putting Small Area Estimation to the Test." Policy Research Working Paper 4155. World Bank, Washington, D.C.

Dibley, M.J., J.B. Goldsby, N.W. Staehling, and F.L. Trowbridge 1987a. "Development of Normalized Curves for the International Growth Reference: Historical and Technical Considerations." *American Journal of Clinical Nutrition* 46(5): 736–48.

Dibley, M.J., N.W. Staehling, P. Nieburg, and F.L. Trowbridge 1987b. "Interpretation of *Z*-Score Anthropometric Indicators Derived from the International Growth Reference." *American Journal of Clinical Nutrition* 46(5): 749–62.

Elbers, C., T. Fujii, P. Lanjouw, B. Özler, and W. Yin 2007. "Poverty Alleviation through Geographic Targeting." *Journal of Development Economics* 83(1): 198–213.

Elbers, C., J.O. Lanjouw, and P. Lanjouw 2002. "Micro-Level Estimation of Welfare." Policy Research Working Paper 2911. World Bank, Washington, D.C.

———. 2003. "Micro-Level Estimation of Poverty and Inequality." *Econometrica* 71(1): 355–64.

Elbers, C., P. Lanjouw, and P.G. Leite 2008. "Brazil within Brazil: Testing the Poverty Map Methdology in Minas Gerais." Policy Research Working Paper 4513. World Bank, Washington, D.C.

Elbers, C., P. Lanjouw, J. Mistiaen, B. Özler, and K. Simler 2003. "Are Neighbours Equal? Estimating Local Inequality in Three Developing Countries." Discussion Paper 147. International Food Policy Research Institute, Food Consumption and Nutrition Division, Washington, D.C.

Foster, J.E., J. Greer, and E. Thorbecke 1984. "A Class of Decomposable Poverty Indices." *Econometrica* 52: 761–66.

Frongillo, E.A., M. de Onis, and K.M.P. Hanson 1997. "Socioeconomic and Demographic Factors Are Associated with Worldwide Patterns of Stunting and Wasting of Children." *Journal of Nutrition* 127(12): 2302–09.

Fujii, T. 2006. “Commune-Level Estimation of Poverty Measure and Its Application in Cambodia.” In R. Kanbur, A.J. Venables, and G. Wan eds, *Spatial Disparities in Human Development: Perspectives from Asia*. Tokyo: United Nations University Press.

———. 2007a. “Geographic Decomposition of Inequality in Health and Wealth: Evidence from Cambodia.” Economics and Statistics Working Paper 24-2007. Singapore Management University, Singapore.

———. 2007b. “To Use or Not To Use? Poverty Mapping in Cambodia.” In T. Bedi, Aline Coudouel, and K. Simler, eds., *More Than a Pretty Picture: Using Poverty Maps to Design Better Policies and Interventions*. Washington, D.C.: World Bank.

———. 2008. “How Well Can We Target Aid with Rapidly Collected Data? Empirical Results for Poverty Mapping from Cambodia.” *World Development* 36(10): 1830–42.

Galler, J.R., and L.R. Barrett 2001. “Children and Famine: Long-Term Impact on Development.” *Ambulatory Child Health* 7: 85–95.

Glewwe, P. 1992. “Targeting Assistance to the Poor: Efficient Allocation of Transfers When Household Income Is Not Observed.” *Journal of Development Economics* 38(2): 297–321.

Glewwe, P., H.G. Jacoby, and E.M. King 2001. “Early Childhood Nutrition and Academic Achievement: A Longitudinal Analysis.” *Journal of Public Economics* 81: 345–68.

Haddad, L., H. Alderman, S. Appleton, L. Song, and Y. Yohannes 2003. “Reducing Child Malnutrition: How Far Does Income Growth Take Us.” *World Bank Economic Review* 17(1): 107–31.

Haughton, D., and J. Haughton 1997. “Explaining Child Nutrition in Vietnam.” *Economic Development and Cultural Change* 45(3): 541–56.

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.” *World Bank Economic Review* 14(1): 147–66.

Kanbur, R. 1987. “Transfers, Targeting and Poverty.” *Economic Policy* 4(1): 112–47.

Li, Y., G. Guo, A. Shi, Y. Li, T. Anme, and H. Ushijima 1999. “Prevalence and Correlates of Malnutrition among Children in Rural Minority Areas of China.” *Pediatrics International* 41: 549–56.

Pelletier, D.L., E.A. Frongillo Jr., , D.G. Schroeder, and J.-P. Habicht 1994. “A Methodology for Estimating the Contribution of Malnutrition to Child Mortality in Developing Countries.” *Journal of Nutrition* 124: 2106S–22S.

Pradhan, M., D.E. Sahn, and S.D. Younger 2003. “Decomposing World Health Inequality.” *Journal of Health Economics* 22(2): 271–93.

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.

Sahn, D.E., and D. Stifel 2003. “Exploring Alternative Measures of Welfare in the Absence of Expenditure Data.” *Review of Income and Wealth* 49(4): 463–89.

Shariff, Z.M., J.T. Bond, and N.E. Johnson 2000. “Nutrition and Educational Achievement of Urban Primary Schoolchildren in Malaysia.” *Asia Pacific Journal of Clinical Nutrition* 4(9): 264–73.

Tarozzi, A., and A. Deaton 2009. “Using Census and Survey Data to Estimate Poverty and Inequality for Small Areas.” *Review of Economics and Statistics* 91(4): 773–92.

Victora, C. 1992. “The Association between Wasting and Stunting: An International Perspective.” *Journal of Nutrition* 122(5): 1105–10.

Waterlow, J.C., R. Buzina, W. Keller, J.M. Lane, M.Z. Nichaman, and J.M. Tanner 1977. “The Presentation and Use of Height and Weight Data for Comparing the Nutritional Status of Groups of Children under the Age of 10 Years.” *Bulletin of the World Health Organization* 55(4): 489–98.

WFP (World Food Programme). 2001. *Identifying Poor Areas in Cambodia: Combining Census and Socio-Economic Survey Data to Trace the Spatial Dimensions of Poverty.* Phnom Penh: World Food Programme.

WHO (World Health Organization). 2002. *World Health Report 2002: Reducing Risks and Promoting Healthy Life*. Geneva: World Health Organization.

WHO (World Health Organization) Working Group. 1986. "Use and Interpretation of Anthropometric Indicators of Nutritional Status." *Bulletin of the World Health Organization* 64(6): 929–41.

———. 1995. "An Evaluation of Infant Growth: The Use and Interpretation of Anthropometry in Infants." *Bulletin of the World Health Organization* 73(2): 165–74.

Wright, J.A., C.A. Ashenburg, and R.C. Whitaker 1994. "Comparison of Methods to Categorize Undernutrition in Children." *Journal of Pediatrics* 124(4): 944–46.

Zeini, L.O., and J.B. Casterline 2002. "Clustering of Malnutrition among Egyptian Children." Cairo University, Cairo, and Population Council, New York.