9-2013

Geographic decomposition of inequality in health and wealth: Evidence from Cambodia

Tomoki FUJII
Singapore Management University, tfujii@smu.edu.sg

Published in Journal of Economic Inequality, September 2013, Volume 11, Issue 3, pp 373–392.
https://doi.org/10.1007/s10888-012-9226-3

Geographic decomposition of inequality in health and wealth: evidence from Cambodia

Tomoki Fujii

Abstract Applying the small-area estimation methods to Cambodia data, we decompose the total inequality in wealth (consumption) and health (child undernutrition) indicators into within-location and between-location components. Because the knowledge of the pattern of spatial disparity in poverty and undernutrition is important for the geographic targeting of resources, we conduct a geographic decomposition of the variance of the Foster-Greere-Thorbecke index in addition to the standard decomposition exercise based on the generalized entropy measures. We find that a sizable proportion of wealth inequality is due to between-location inequality, whereas health inequality is mainly due to within-location inequality.

Keywords Small-area estimation · Health inequality · Cambodia

1 Introduction

It is widely known that health and wealth are positively correlated. However, it is not clear whether the spatial inequalities in health and wealth necessarily exhibit a similar pattern. In this paper, we empirically investigate the importance of spatial inequality in health and wealth, for which we use (transformed) z-scores and per capita consumption measures as a proxy. We find that the relative importance of

Earlier stages of this research were supported by the World Bank Research Group, the Government of Japan under the Millennium PHRD grant, and the SMU Research Grant (05-C208-SMU-030). I thank Alain de Janvry, Chris Elbers, Jean Olson Lanjouw, Peter Lanjouw, Livia Montana, Mahadevan Ramachandran, Martin Ravallion, H.E. Kim Saysamalen, Elisabeth Sadoulet, Boreak Sik, H.E. San Sythan, and two anonymous referees for comments and help I received at various stages of my research. I also benefited from the discussion with the participants of the APL seminar at the IDE-JETRO.

T. Fujii (✉)
School of Economics, Singapore Management University, 90 Stamford Road, Singapore 178903, Singapore
e-mail: tfujii@smu.edu.sg

between-location inequality to total inequality is much higher for wealth inequality than for health inequality; moreover, this finding is robust to the choice of inequality index and transformation of outcome variables.

It is useful to understand the extent to which inequality exists within and between locations for the following reason: policy-makers may want to target assistance programs to the less healthy or wealthy within a population. The effectiveness of geographic targeting of such programs depends on the extent to which the less healthy or wealthy are spatially concentrated. We, therefore, investigate the contribution of between-location inequality to the total inequality.

The following simple example further elucidates this point: Suppose that wealthy people in a particular country tend to live in the north and poor people tend to live in the south. Suppose further that mosquitoes carrying malaria parasites exist uniformly across the country. Because wealthy people have better knowledge of how to cope with malaria and resources to prevent the infection (such as mosquito repellants and mosquito nets), they are less likely to become infected compared to poor people. However, as there is no perfect preventive measure, the incidence of malaria would be less unequally distributed across the country than poverty.

Knowledge of the pattern of spatial inequalities in health and wealth is valuable for geographic targeting because the spatial inequalities characterize the potential gains from geographic targeting. In the example given above, the resources for anti-poverty programs can be most efficiently used if they are delivered to the south because everyone there is poor; hence, the resources all go to poor people. However, if we deliver all of the resources for health programs (say, malaria tablets) to the south, the outcome may not be fully efficient. We may be providing malaria tablets to those in the south who are not vulnerable to malaria while not providing tablets to those in the north who are vulnerable.

If geographic information is the only information available to the policy-maker, geographic targeting is still useful (and efficient given the available information), but the extent to which one may gain from geographic targeting is determined by the pattern of spatial inequality. Therefore, it is useful to understand how different the spatial patterns of inequality of health and wealth are. As elaborated later, the spatial patterns of inequality that focuses on the lower tail of the distribution (i.e., poverty and undernutrition as opposed to consumption and transformed z-scores) are particularly relevant.

While the example discussed above is extreme, it is nonetheless relevant to the situation in Cambodia. On the one hand, the mean consumption expenditure at the provincial level varies quite substantially across Cambodia. Based on the Cambodia Socio-Economic Survey (CSES) for 1997, the maximum and minimum provincial-level consumption expenditures per day per capita, in terms of purchasing power parity (PPP), were USD 6.33 and USD 0.78, respectively.[1] On the other hand, the mean height-for-age z-score for children under the age of five—an indicator of overall health status that will be discussed in detail later—does not significantly vary across the country, according to the Cambodia Demographic and Health

[1] The provincial-level estimates may not be very accurate because the CSES data are only representative for three strata: Phnom Penh, Other Urban, and Rural [22]. The PPP USD figures are calculated using the private consumption conversion factor of USD 1 = KHR 1445 for 1997, taken from the World Development Indicators.

Survey (CDHS) for 2000.[2] The maximum and minimum at a comparable level of aggregation[3] are −1.2 and −2.2, respectively. Child undernutrition is widespread, even in relatively wealthy provinces, because child care is often inadequate and diseases such as malaria and diarrhea are prevalent.

This difference in the spatial disparities between the health and wealth indicators implies that geographic targeting may be effective for anti-poverty policies but less so for policies designed to improve the nutritional status of children. It is, therefore, inappropriate to conclude that delivering food aid to the poor would automatically reduce child undernutrition in Cambodia.

One may object to this argument for two reasons. First and most importantly, it could be argued that the measures we use above are not directly comparable because the measurement levels may not be the same. We address the first point by using categorical (binary) outcome variables, which do not require that the outcome variables be measurable on a ratio scale.

The second possible objection is that we are only looking at the minimum and maximum at one level of aggregation. Therefore, it is not clear whether our argument holds if we look at other statistics or different levels of aggregation. We deal with this point by applying the small-area estimation method, in which a survey and a census are combined to derive the estimates of health and wealth statistics for small areas. We then decompose the total inequality into between-location inequality and within-location inequality at various levels of aggregation. Even if we use a categorical measure and decompose the data at a much more disaggregated level than provinces, our main conclusion still holds: in Cambodia, health inequality is primarily an individual phenomenon, whereas a sizable portion of wealth inequality can be explained by spatial disparities.

Our decomposition exercise allows us to pin down the levels of aggregation at which a substantial amount of spatial disparity is observed. The decomposition results are also informative for practical purposes. They help us determine the level of aggregation at which targeting should be performed. For example, if the spatial inequality is largely due to between-province inequality, it may be best to conduct a few large projects in a small number of provinces. Conversely, if the spatial inequality is largely due to between-village inequality in a given commune, it may be best to carry out a large number of small projects in many villages. As Kanbur [19] emphasizes, the decomposition exercise alone cannot determine what policy instruments are most desirable because the costs and magnitudes of the impacts of the specific policy instruments considered must be taken into account. Nevertheless, our decomposition exercise provides valuable information that could be used to determine the types of policies and the target aggregation levels.

This paper is organized as follows. In the next section, we review the related literature on the relationship between health and wealth. We provide a discussion of the data and measurements in Section 3, which is followed by a discussion of the small-area estimation methods used in this study in Section 4. In Section 5, we shall

[2]The z-score measures used in this study are based on the NCHS/WHO international reference population to maintain consistency with the official figures published by the National Institute of Statistics, Directorate General for Health, and ORC Macro [24].

[3]To be more precise, the calculation is based on seventeen strata. Thirteen of these strata consist of a single province, whereas the remaining four are combinations of multiple provinces.

discuss the inequality decomposition method. We then present the decomposition results in Section 6. Section 7 provides some discussion.

2 Review of related studies

This study is related to the literature on the positive relationship between health and wealth. Broadly speaking, there are three types of explanations for this relationship: (i) health causes wealth, (ii) wealth causes health, and (iii) there is a tertiary factor that is correlated with both health and wealth. Let us briefly look at each type of explanation.

Prichett and Summers [26] confirm that wealth causes health, using instrumental variables to isolate reverse causation or incidental associations. The long-run income elasticity of infant and child mortality in developing countries are estimated at between -0.2 and -0.4. The results are intuitive, as one would expect that wealthier populations tend to enjoy more and better health care, food, and water.

Bloom and Canning [1], however, point out several mechanisms that could account for this relationship. First, healthier populations tend to have higher productivity. Second, a healthier population has stronger incentives to save and invest in human and physical capital. Third, demographic dividends are also a possible mechanism. That is, when the decline in infant mortality initiates a decline in fertility, the proportion of the working age population increases. As a result, per capita income in the country also increases.

The positive association between health and wealth can also be explained by the existence of a third factor that is correlated with both. For example, a better educated population may well have higher incomes and better information about health that helps them to remain healthy. Indeed, Elo and Preston [11] show that a one-year increase in education for persons aged 35–54 is associated with a reduction in mortality of approximately 7–9% for males and 2–8% for females in various wealthy countries. In addition to education, health care and behavior may also affect both health and wealth.

The three types of relationships mentioned above are not mutually exclusive and may be at work simultaneously. They have, however, very different policy implications. For example, suppose that education is the most important determinant of health. In this case, income redistribution may not help to reduce health inequality. Reducing health inequality would require, for example, providing an equal educational opportunity to everyone.

Another important and controversial point in the relevant literature is whether absolute or relative income is what matters in determining health outcomes. Some researchers, including Wilkinson [35–37] and Kawachi et al. [20], have noted that poor health in developed countries is strongly related to income inequality. Kawachi et al. [20] argue that income inequality is related to a reduction in social cohesion, which in turn is associated with poor health (as measured by mortality). Marmot and Wilkinson [21] point out that the psychosocial effects of relative deprivation, including control, anxiety, insecurity, depression, and social afflictions, are negatively correlated with health. These studies would suggest that inequalities in health outcomes may be caused by income inequality. This, in turn, provides the rationale for income redistribution as a means to reduce health inequality.

Economists have been skeptical about this argument. Wagstaff and van Doorslaer [31] point out that the empirical evidence in earlier studies based on aggregate data is often insufficient to discriminate one hypothesis from another. Based on the analysis of individual-level data in the United States, they found strong support for the absolute-income hypothesis (health is determined by the level of income) and little or no support for the relative-income hypothesis (health is determined by the deviation in income from the mean) or the income-inequality hypothesis (health is determined by the level of income and income inequality). Deaton [2] provides a more comprehensive critique of the relative-income and income-inequality hypotheses.

This study also investigates the relationship between health and wealth, but we look at a different aspect, namely, spatial inequalities in health and wealth. Few studies have been conducted addressing spatial inequality. Using a cross-country regression analysis, van Doorslaer et al. [29] directly relate a health inequality index to an income inequality index and find that the estimated coefficient is positive and significant. Pradhan et al. [25] decomposed world health inequality into within-country inequality and between-country inequality; the latter accounts for only 31% of total world inequality. This is in contrast to similar exercises for income inequality; Various empirical studies agree that between-country inequality accounts for more of the total inequality than within-country inequality [12]. This contrast is consistent with our finding that the relative importance of between-location inequality to within-location inequality is much larger for wealth inequality than for health inequality.

While these cross-country comparisons are interesting, an inequality decomposition within a country is arguably more important because it is more closely related to the geographic targeting of health programs, anti-poverty programs, or a combination of both. In this sense, our study is closely related to that of Wagstaff [30], who asks, "How far are income-related inequalities in the health sector due to gaps *between* poor and less poor areas, rather than due to differences between poor and less poor people *within* areas." He lays out a method to answer this question and applies it to a geographical decomposition of concentration indices of health subsidies in Vietnam and insurance coverage in rural China.

This study differs from Wagstaff [30] in three respects. First, instead of "health access" type indicators, we use individual health outcomes. Second, by applying small-area estimation, we are able to decompose at various levels of decomposition, covering virtually all of the country. In contrast, the number of geographic areas in Wagstaff [30] was relatively small (58 provinces in Vietnam and 225 villages in China) and the surveys used in his study cover only a fraction of each geographic area.[4] Third, Wagstaff [30] uses the concentration index, which can only be decomposed into the between-area component, within-area components, and the residual. The residual can be interpreted, but the interpretation is not straightforward, a feature that makes the concentration index unattractive. Hence, we use various decomposable measures that allow for a neat decomposition, as we shall discuss in Section 6.

[4]Furthermore, the surveys may not be representative at the level of decomposition. For example, the 1998 Vietnam Living Standard Survey used in Wagstaff [30] is not representative at the provincial level. Hence, the within-province concentration index is unreliable at best.

Finally, this study is related to Elbers et al. [10] in that both compare the shares of between-group inequality in the total inequality, taken from multiple sources. Elbers et al. [10] compare the proportion of between-group inequality among different social groups in eight countries. They note that each country has a different number of social groups that are used to calculate between-group inequality. They argue that the proportion of between-group inequality calculated in the standard way may not permit appropriate comparisons because between-group inequality is likely to be larger when the number of social groups is larger. Therefore, they propose a calculation of the ratio of the observed between-group inequality to the maximum between-group inequality that can be generated from the observed distribution of consumption and observed size of each social group.

This study also attempts to compare inequality decomposition results. Unlike Elbers et al. [10], we compare the inequality decomposition results for three different outcome indicators in Cambodia. Because the definition of groups used for inequality decomposition is the same across the three indicators, we do not face the same problem as Elbers et al. [10].

3 Data and measurement

There are three primary data sources used in this study. The first source is the CSES 1997. It covers 6,010 households from 474 villages, and villages were used as the primary sampling unit. It is representative at three strata: Phnom Penh, Other Urban, and Rural. The CSES data include a consumption indicator and various other indicators, both at the household and individual levels. We use the logarithmic consumption per capita in the household as a proxy for the wealth level (Further details on CSES 1997 are given in [22]).

The second source is the CDHS 2000, which covers health and demographic information for the Cambodian population with a focus on women of childbearing age and young children. The sample covered 12,236 households in 17 strata across the country. In addition to detailed information about each household, its members, and housing characteristics, one-quarter of these households were randomly selected to participate in the collection of anthropometric data. All children under 60 months of age in the sub-sampled households were weighed and measured. After excluding children for which information on height or weight is missing or implausible, 3,596 observations were used for this analysis (for details, see [24]).

As indicators of health, we use the standardized height and weight for children under the age of five, which are essentially the height-for-age and weight-for-age z-scores but expressed in terms of the height and weight of a child with the same z-scores in the reference population of healthy children.[5] The concept of standardized height was first proposed by Pradhan et al. [25]. They adopted this concept to avoid negative numbers, which was necessary for inequality decomposition. Fujii [18] also adopted this measure and constructed a similar measure for weight. In general, excessive weight is unhealthy, and thus, standardized weight may not appear to be an

[5] We take the base reference population to be 24-month-old girls, following Pradhan et al. [25]. We check the robustness of our results with respect to the choice of reference population.

appropriate measure of health. However, this is not an important point in Cambodia because the proportion of overweight children is almost negligible.

Three remarks are in order. First, the height-for-age and weight-for-age z-score measures are widely accepted indicators of the nutritional status of children. A lack of nutrition is closely related to health because undernutrition is associated with mortality and morbidity in later life as well as various mental, cognitive, and behavioral problems. Therefore, the z-score measures are also reasonable proxy measures for overall health status, even though some important dimensions of health may not be appropriately captured by these measures. Second, as with the z-scores, the standardized height and weight are not direct measures of nutrition. Therefore, they are unable to distinguish, for example, normal shortness and shortness due to undernutrition. It is, however, still possible to interpret the z-score measures as proxies for health status because they reflect the risk of the adverse effects of undernutrition [38].

Finally, while both the standardized height and weight reflect the child's nutritional status, they are different indicators with different interpretations [33]. A low height-for-age z-score below the conventional cut-off of negative two is called stunting. Stunting reflects chronic, or long-term, undernutrition because one cannot lose height and a catch-up in height takes a very long time.

Conversely, a low weight-for-height z-score, or wasting, reflects acute, or short-term, undernutrition. This is because people lose weight when nutrition is compromised, but they cannot lose height. A low weight-for-age z-score, or underweight, occurs as a result of both recent nutritional shortfalls and accumulated stunting.[6]

The third data source is the Cambodian National Population Census (CNPC) for 1998 (for details, see National Institute of Statistics [23]). The census covers almost all persons residing in Cambodia at the time of the census and includes variables for housing characteristics, conditions, and facilities as well as individual variables for sex, age, relation to the head of household, marital status, migration, literacy, education, and employment. After excluding the records with missing values, the census contains approximately 2.1 million household records and 1.4 million records of children under five.

In addition to these data sources, we used various data sources that could be merged into the census and the survey. These sources cover indicators such as distance calculations, land use and land cover information, climate indicators, vegetation, agricultural production, and flooding as well as the village-level means generated from the census. The inclusion of these geographic variables and their cross terms with other individual-level and household-level variables has substantially improved the ability to explain the variation of consumption and anthropometric indicators.

4 Small-area estimation

As mentioned earlier, we apply small-area estimation methods to decompose inequality measures at geographically disaggregated levels. The basic idea of the

[6]See Dibley et al. [4, 5], Waterlow et al. [32], and WHO Working Group [33, 34] for further discussions of the z-score measures.

small-area estimation is simple; An econometric model is estimated with a survey, which contains the outcome variables of interest. Then, the outcome variables are imputed into the census, which covers the entire population of interest. The imputed variables can be used, among other things, to produce poverty measures for small areas or to perform the decomposition analysis described in the next section.

This idea was first studied rigorously by Elbers et al. [8, 9]. Their method has been applied to obtain poverty estimates for small areas in many developing countries, and it has become one of the standard tools for poverty analysis. The small-area estimates of poverty measures are typically plotted on a map, which is commonly called a poverty map. Poverty maps have proved useful for the analysis and formulation of anti-poverty policies, and they have become increasingly popular among policy-makers and researchers.

In addition to creating poverty maps, small-area estimation has been used for a wide array of purposes. For example, it has been used to analyze geographic targeting [6, 16], consumption inequality [7], local inequality and crime [3], and the impacts of trade liberalization [17]. This paper offers a new application of small-area estimation; we use small-area estimation to compare the importance of spatial inequalities in health and wealth through a decomposition analysis.

Fujii [14] produced poverty maps for Cambodia by combining the CSES 1997 and CNPC 1998. These poverty maps have been used by various international organizations, ministries, and non-governmental organizations to analyze the poverty situations in their operational areas, selecting potential locations for their projects and programs, and educating students in classrooms [15].

Fujii [18] extended the small-area estimation method to estimate the nutritional status of children and applied his method to Cambodia. In his study, the CDHS 2000 and CNPC 1998 are combined to create nutrition maps in Cambodia, which give the estimates of the prevalence of undernutrition. The decomposition analysis carried out in this paper uses inequality estimates that can be obtained as a by-product of the small-area estimation methods. In the remainder of this section, we shall discuss the basic steps of these methods. Further details on small-area methods and their applications to Cambodia are found in Elbers et al. [8, 9] and Fujii [14, 18].

In small-area estimation methods, the following model is typically considered:

$$y_l = E[y_l\,|\mathbf{x}_l\,] + u_l = \mathbf{x}_l \cdot \beta + u_l,$$

where y_l, $\mathbf{x}_l$, and u_l are, respectively, the outcome indicator of interest, the vector of regressors, and the unobserved random error term for the observation unit l. Note that $\mathbf{x}_l$ must be indicators that are included both in the survey and census.

To conduct of the decomposition analysis for wealth inequality, we use logarithmic consumption per capita as the outcome indicator of interest. The observation unit in this case is a household. We call this application of small-area estimation poverty mapping. For the decomposition analysis of health inequality, we use the standardized height or weight as the outcome indicators of interest. In this case, the observation unit is an individual. We refer to this application of small-area estimation as nutrition mapping.

We run a generalized-least-squares regression to obtain the point estimate $\hat{\beta}$ of the model coefficient β and its associated variance matrix from the survey. We also obtain an estimate of the distribution of u_l, which may have subcomponents, such as cluster-specific and household-specific random errors.

We then carry out Monte Carlo simulations to explicitly account for the model error (estimation error of $\hat{\beta}$) and the idiosyncratic error (error due to u_l). In each round of the Monte-Carlo simulation, we first draw the model coefficient. Further, for each observation unit, we randomly draw the error term u_l in a manner consistent with the estimated distribution of u_l. Thus, we have an imputed outcome indicator for each observation unit l in each round of the Monte-Carlo simulation, which in turn can be transformed and aggregated at any level to arrive at the "welfare" indicator of interest, such as the FGT measure of poverty due to Foster et al. [13]. By taking the mean and standard deviation covering all of the rounds of the simulation, we obtain the point estimates and their standard errors for the welfare indicator. The point estimates are typically presented in the form of a map.

Although the basic principles of small-area estimation are the same between Elbers et al. [8, 9] and Fujii [18], there are four major differences. First, the outcome indicator is different. Second, the structures of the random error terms are different. Because the anthropometric indicators are measured at the individual level, Fujii [18] allows for individual-specific random effects in addition to the cluster-specific and household-specific random effects included in Elbers et al. [8, 9]. Third, because there are multiple health outcome indicators (i.e., standardized height and weight), the correlations of the error terms across different outcome indicators are allowed in Fujii [18]. However, this is not an issue in Elbers et al. [8, 9] because they only consider a single outcome indicator. Finally, Fujii [18] makes finite-sample corrections to account for, among other things, the fact that there are only a few children under five in a typical household. Elbers et al. [8, 9] do not make such corrections.

5 Decomposition analysis

Because we are interested in the shares of between-location inequality in health and wealth, we use the generalized entropy (GE) measure of inequality, which is additively decomposable and satisfies the desirable characteristics of inequality measures, such as the transfer principle, scale independence, population-replication independence, and anonymity [27, 28]. To introduce the generalized entropy measure, let us define by w_l the weight attached to each observation unit. In our application, this weight is the household size for poverty mapping and unity for nutrition mapping. Furthermore, let us define the following:

$$f_\alpha(y_l, \bar{y}) \equiv \begin{cases} \ln(\bar{y}/y_l) & \text{if} \quad \alpha = 0 \\ y_l/\bar{y} \cdot \ln(y_l/\bar{y}) & \text{if} \quad \alpha = 1 \\ \left[(y_l/\bar{y})^\alpha - 1\right] / [\alpha(\alpha - 1)] & \text{if } \alpha \neq 0, \alpha \neq 1 \end{cases}.$$

The generalized entropy measure with parameter α for group g can then be written as follows:

$$GE_g(\alpha) \equiv W_g^{-1} \sum_{l \in g} w_l f_\alpha\left(y_l, \bar{y}_g\right),$$

where $W_g \equiv \Sigma_{l \in g} w_l$ and $\bar{y}_g \equiv W_g^{-1} \Sigma_{l \in g} w_l y_l$ are the sum of weights and the weighted average of the outcome indicator for group g, respectively. Pradhan et al. [25] set the parameter value at $\alpha = 0$, which means that everyone is equally weighted. While this choice is sensible, as there is no other obvious choice, we varied the parameter values to see if our results are sensitive to the choice of the parameter.

To conduct the inequality decomposition, let C be the set of observation units in a country. Suppose that the country consists of J areas and each observation unit belongs to one and only one of these areas. We denote the set of observation units in the j-th area by G_j. It is well known that the total inequality GE_C for a country can be written as the sum of the between-group component GE_C^B and the within-group component GE_C^W, where these components are defined as follows:

$$GE_C^B(\alpha) = W_C^{-1} \sum_{j=1}^{J} W_{G_j} f_\alpha\left(\bar{y}_{G_j}, \bar{y}_C\right) \text{ and}$$

$$GE_C^W(\alpha) = W_C^{-1} \sum_{j=1}^{J} W_{G_j}\left(\bar{y}_{G_j}/\bar{y}\right)^\alpha GE_{G_j}(\alpha).$$

Thus, the within-group component is the weighted average of the inequality for each group, and the between-group component is the inequality of group-level means. Using the notations introduced above, we can calculate the share of between-group inequality by $GE_C^B(\alpha)/GE_C(\alpha)$.

When we calculate health inequality, or inequality in standardized height and weight, we account for the inequality resulting from genetic variation, as did Pradhan et al. [25]. That is, the standard height and weight are intrinsically unequal, and some level of inequality, which we call the natural inequality ($GE_{Natural}$), exists in a healthy population. This means that the natural inequality exists within each area even when the entire population is healthy.

To address this issue, Pradhan et al. [25] adjusted the denominator (the total inequality in the world) by subtracting the natural inequality when calculating the proportion of the between-country inequality in total world inequality, where the natural inequality is calculated as the inequality of standardized height for the reference population of healthy children. Put differently, they augmented the proportion of between-country inequality in total world inequality by the correction factor $f = GE_{World}(0)/\left[GE_{World}(0) - GE_{Natural}(0)\right]$.

Because the same issue exists in our case, we adjust the proportion of between-group inequality in a similar manner. We subtract the natural inequality from the denominator (i.e., the total inequality in Cambodia) when calculating the share of between-location inequality. Thus, the correction factor we use to augment the proportion of between-location inequality in total inequality in Cambodia is $f = GE_{Cambodia}(\alpha)/\left[GE_{Cambodia}(\alpha) - GE_{Natural}(\alpha)\right]$ for all α in this study. Notice that this adjustment unambiguously increases the proportion of between-location inequality.

While the share of between-location inequality is an interesting descriptive statistic, we may be more interested in the spatial inequality of certain types of populations. This is because standardized height and weight are not measures for which increasing values represent welfare improvements across the entire distribution. For example, increasing height by half a standard deviation for someone who is three standard deviations below the mean of the reference population is clearly an improvement, but the same cannot be said about someone at the mean. Therefore, when we are interested in the spatial concentration of the poor or undernourished, statistics that focus on the lower tail of the distribution are more appropriate.

This point is also important from the perspective of targeting. Because we can simultaneously have low or zero between-group inequality and heterogeneous levels of undernourishment or poverty across areas, the share of between-group inequality may not be very informative regarding the potential usefulness of geographic targeting.

This point may be more clearly understood with a simple numerical example. Suppose that there are two villages, A and B, in a small country, and each village has 100 people. In Village A, everyone earns 10. In Village B, there is one rich person whose income is 901, and the remaining 99 people earn only 1. Let us suppose that the poverty line is 5. In this case, the between-group inequality measured by the generalized entropy measures is zero because the average income is 10 in both villages. However, geographic targeting is obviously very useful because all the poor people live in Village B. This shows that what really matters for the targeting of poverty alleviation programs is not the inequality of income but the "inequality of poverty" across areas. A similar argument holds for the targeting of child nutrition programs.

To address this issue, we conduct a decomposition analysis that focuses on undernutrition and poverty. To this end, we consider the FGT measure for each observation unit, which is defined as follows:

$$P_l(\alpha) = \left(\frac{z - y_l}{z}\right)^{\alpha} \operatorname{Ind}(y_l < z),$$

where z is the threshold below which the observation unit l is deemed poor or undernourished.[7] The FGT measure for group g is simply $\bar{P}_g(\alpha) \equiv W_g^{-1} \sum_{l \in g} w_l P_l(\alpha)$. With these definitions, we can define the variance of the FGT measure for group g as follows:

$$V_g(\alpha) \equiv W_g^{-1} \sum_{l \in g} w_l h_\alpha\left(P_l, \bar{P}_g\right) \text{ where } h_\alpha\left(P_l, \bar{P}_g\right) \equiv \left(P_l(\alpha) - \bar{P}_g(\alpha)\right)^2.$$

It is straightforward to show that we can decompose the total variance $V_C(\alpha)$ of the FGT measure in country C into the within-group component $V_C^W(\alpha)$ and between-group component $V_C^B(\alpha)$ in the following manner:

$$V_C(\alpha) = W_C^{-1} \sum_{j=1}^{J} W_{G_j} V_{G_j}(\alpha) + W_C^{-1} \sum_{j=1}^{J} W_{G_j} h_\alpha\left(\bar{P}_{G_j}, \bar{P}\right) \equiv V_C^W(\alpha) + V_C^B(\alpha).$$

As with the decomposition of generalized entropy measures, the within-group component is the weighted average of the variance for each group, and the between-group component is the variance of the group-level means. We take the proportion of the between-group variance to be V_C^B / V_C, where the denominator is adjusted for natural inequality in the case of standardized height and weight. By comparing this

[7] We have chosen thresholds that ensure comparability with previous studies. For the poverty line, we use one that is consistent with the official poverty estimates [14]. For the cutoff of undernutrition, we use the standardized height and weight corresponding to a z-score of negative two.

proportion across different outcome variables, we can tell whether poverty is more or less spatially concentrated than undernutrition in Cambodia

Three remarks are in order. First, we can readily compare the spatial inequality in poverty and undernutrition when $\alpha = 0$. This is because the FGT measures at the unit-record levels are simply binary indicators in this case, and thus, a comparison does not require the multiplicability of the outcome variable. Therefore, once the definitions of poverty and undernutrition are accepted, no comparability issue arises for the variance decomposition of the $P(0)$ index.

Second, the proportion of the between-group variance does not depend on the (arbitrary) choice of the reference population when $\alpha = 0$. This is because under-nourishment is determined by whether the child's z-score is below the threshold and thus it does not depend on the reference group. Note that the two points made above are not valid for $\alpha \neq 0$. For example, there is no clear comparison between the poverty gap and the undernutrition gap (i.e., how far below y is from z as a proportion of z) when $\alpha = 1$ because the undernutrition gap in part depends on the choice of the reference group.

Third, this decomposition analysis ignores what happens in the upper tail of the distribution because the FGT measures are not sensitive to the distribution above the threshold. This asymmetry between the lower and upper tails is important for the purposes of policy analysis and formulation. For those implementing anti-poverty programs, a transfer of income from the poorest to the second poorest would be problematic, as such a transfer would increase $P(\alpha)$ for $\alpha > 1$. However, a small transfer from the second richest to the richest would not be problematic. Hence, the variance decomposition method presented above is potentially more relevant to policy-makers than the decomposition of the generalized entropy measures.

6 Results

As the discussion in Section 3 makes clear, we obtain an imputation of the outcome variable for each observation unit in each round of the simulation. Therefore, we can compute not only the usual "welfare" indicators such as the FGT measures of poverty but additional measures such as the generalized entropy measures. To see the spatial pattern of inequality, we have plotted the commune-level estimates of $GE(0)$ for consumption, standardized height, and standardized weight measures in Figs. 1, 2, and 3, respectively; a commune is the third largest administrative unit after the province and district and before the village.[8] Because the comparison of $GE(0)$ in absolute terms is not meaningful, we simply looked at the quartiles. The lighter areas on the maps represent more equal communes. For example, Q1 is the top quarter of the most equal communes.

There are four points worth noting. First, these maps allow us to identify, for example, the areas in which the level of local inequality, or within-location inequality, in health or wealth is high. Hence, if policy-makers are concerned about local inequality, these maps are helpful. Otherwise, however, these maps may not be particularly informative for targeting purposes because they do not tell us the

[8]There are about 1,600 communes in Cambodia and each commune contains around 1,300 households.

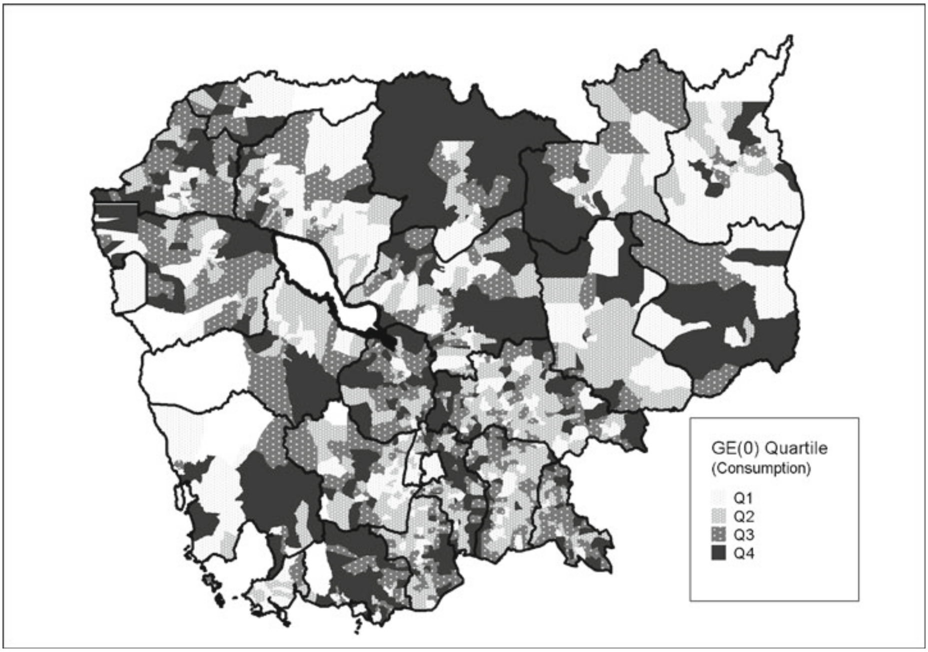

Fig. 1 Map of inequality (GE(0)) in consumption at the commune level

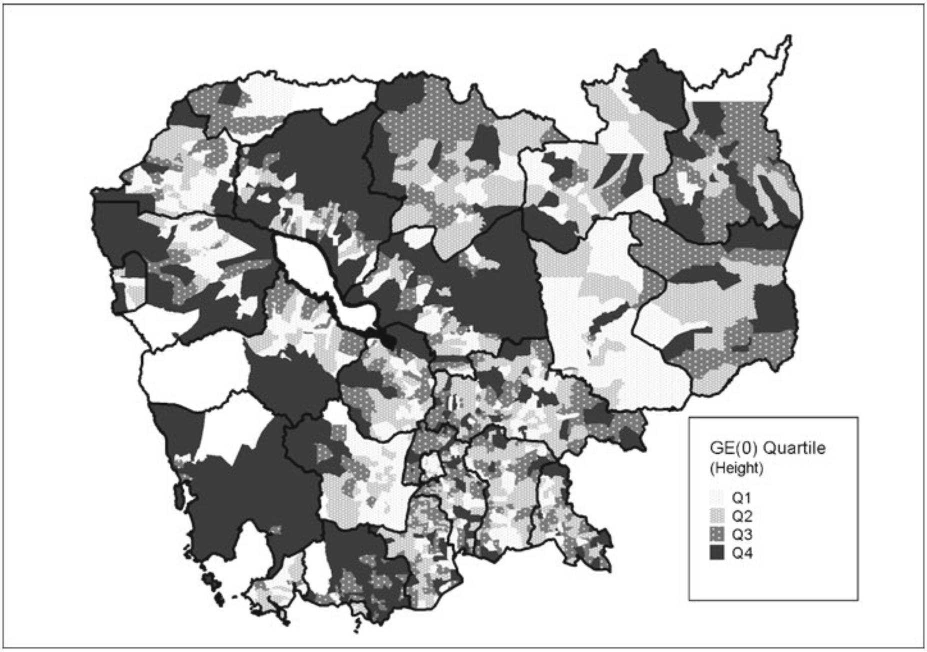

Fig. 2 Map of inequality (GE(0)) in standardized height at the commune level

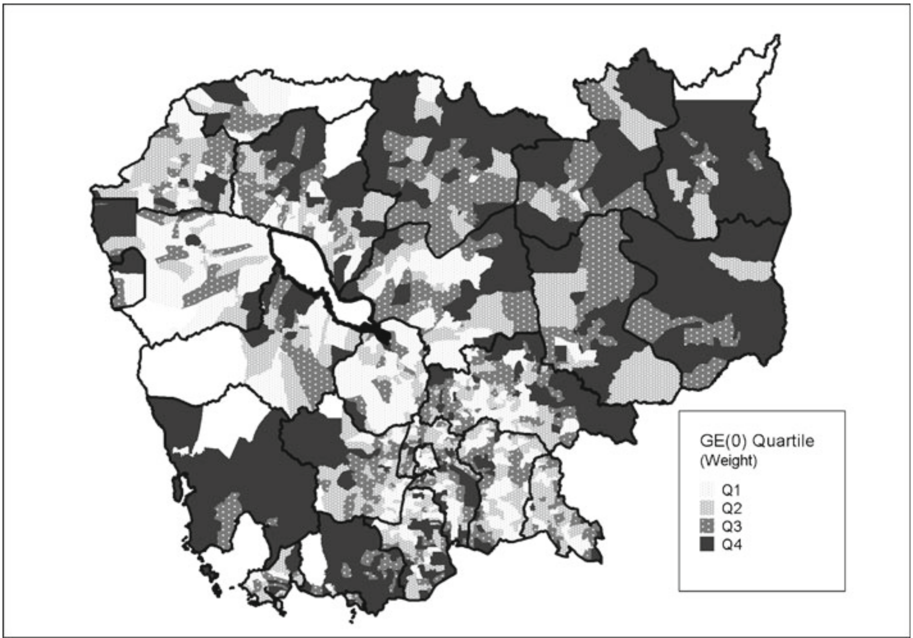

Fig. 3 Map of inequality (GE(0)) in standardized weight at the commune level

importance of between-location inequality relative to total inequality. Second, the size of a commune is very heterogeneous across Cambodia. In general, the area of a commune is larger in more remote and sparsely populated areas such as the northeastern part of the country. This means that remote communes tend to be overrepresented in the map. Therefore, the users of the maps should not be misled by the apparent size of these remote communes.

Third, these maps do not take into account the standard errors associated with the estimates. Thus, even if we choose one commune from one quartile and another commune from another quartile, the difference in inequality may not be statistically significant. Finally, as $GE(0)$ is sensitive to the lower tail of the distribution, it is responsive to the worst-off children or households.

With these points in mind, let us look at the maps. One can see that the maps for consumption and standardized height are somewhat similar. Indeed, there is a moderately positive correlation between consumption inequality and height inequality; their Spearman rank-correlation is 0.20. However, the rank-correlations of weight inequality with consumption or height inequality are much weaker and negative (−0.01 for consumption and −0.02 for height). This may be because weight is more responsive to short-run negative shocks than height and consumption.

Let us now look at the decomposition results shown in Tables 1 and 2. The figures for the generalized entropy measures are reported in Table 1. The decomposition of the variance of the FGT measures is reported in Table 2. We calculate the proportion of the between-group component in each round of the simulation and take the mean and standard deviation over the simulations to find the point estimate and the

Table 1 Proportion of the between-group component for the generalized entropy measures at various levels[a]

		Cambodia		Province		District		Commune		Village		Unit record	
Standardized height	GE(−1)	0.00	(0.00)	2.28	(0.57)	5.22	(0.68)	13.87	(1.34)	25.01	(1.77)	100.00	(0.00)
	GE(−0.5)	0.00	(0.00)	2.26	(0.57)	5.16	(0.67)	13.76	(1.34)	24.80	(1.76)	100.00	(0.00)
	GE(0)	0.00	(0.00)	2.21	(0.55)	5.04	(0.65)	13.49	(1.32)	24.30	(1.73)	100.00	(0.00)
	GE(0.5)	0.00	(0.00)	2.07	(0.52)	4.71	(0.61)	12.65	(1.25)	22.78	(1.63)	100.00	(0.00)
	GE(1)	0.00	(0.00)	2.56	(0.64)	5.82	(0.75)	15.70	(1.57)	28.27	(2.04)	100.00	(0.00)
Standardized weight	GE(−1)	0.00	(0.00)	0.98	(0.27)	2.61	(0.40)	6.64	(0.76)	14.93	(1.39)	100.00	(0.00)
	GE(−0.5)	0.00	(0.00)	1.00	(0.28)	2.66	(0.41)	6.76	(0.77)	15.18	(1.41)	100.00	(0.00)
	GE(0)	0.00	(0.00)	1.01	(0.28)	2.69	(0.41)	6.84	(0.77)	15.35	(1.41)	100.00	(0.00)
	GE(0.5)	0.00	(0.00)	1.01	(0.28)	2.67	(0.41)	6.81	(0.77)	15.29	(1.40)	100.00	(0.00)
	GE(1)	0.00	(0.00)	1.07	(0.30)	2.83	(0.44)	7.22	(0.81)	16.21	(1.48)	100.00	(0.00)
Consumption	GE(−1)	0.00	(0.00)	15.20	(11.07)	24.22	(10.47)	36.56	(9.26)	57.87	(7.22)	100.00	(0.00)
	GE(−0.5)	0.00	(0.00)	17.64	(14.30)	27.49	(13.48)	41.01	(11.33)	62.17	(7.55)	100.00	(0.00)
	GE(0)	0.00	(0.00)	18.69	(16.32)	28.84	(15.84)	43.19	(13.16)	64.14	(8.68)	100.00	(0.00)
	GE(0.5)	0.00	(0.00)	17.04	(12.66)	27.18	(14.47)	42.25	(12.99)	63.66	(9.09)	100.00	(0.00)
	GE(1)	0.00	(0.00)	11.71	(2.59)	19.91	(3.54)	34.42	(4.06)	57.24	(6.12)	100.00	(0.00)
Number of observations		1		24		180		1,594		13,320/13,233		1,424,907/2,130,544	

[a] The number of villages for standardized height and weight are 13,320, whereas the number for consumption is 13,233. This discrepancy arises because there are some villages without children under five. The figures for standardized height and weight are adjusted for the natural inequality except for Unit Record. The number of unit records corresponds to the number of individuals for standardized height and weight (1,424,907), whereas it corresponds to the number of households for consumption (2.130.544)

Table 2 Proportion of the between-group component in the variance of FGT measures at various levels[a]

		Cambodia		Province		District		Commune		Village		Unit record	
Standardized height	P(0)	0.00	(0.00)	1.34	(0.32)	2.81	(0.37)	6.37	(0.47)	12.66	(0.81)	100.00	(0.00)
	P(1)	0.00	(0.00)	1.28	(0.37)	3.16	(0.50)	7.96	(1.03)	14.97	(1.32)	100.00	(0.00)
	P(2)	0.00	(0.00)	0.97	(0.32)	2.57	(0.49)	7.12	(1.39)	12.87	(1.61)	100.00	(0.00)
Standardized weight	P(0)	0.00	(0.00)	1.04	(0.29)	2.34	(0.33)	5.31	(0.50)	12.01	(0.92)	100.00	(0.00)
	P(1)	0.00	(0.00)	2.16	(0.68)	3.34	(0.72)	5.92	(0.92)	11.60	(1.15)	100.00	(0.00)
	P(2)	0.00	(0.00)	2.14	(0.68)	3.00	(0.73)	4.84	(0.94)	8.97	(1.07)	100.00	(0.00)
Consumption	P(0)	0.00	(0.00)	7.10	(0.50)	12.00	(0.71)	19.06	(0.88)	37.14	(1.16)	100.00	(0.00)
	P(1)	0.00	(0.00)	8.20	(0.86)	14.29	(1.36)	23.15	(1.76)	46.72	(2.01)	100.00	(0.00)
	P(2)	0.00	(0.00)	7.20	(1.00)	12.98	(1.81)	21.53	(2.36)	45.87	(2.82)	100.00	(0.00)
Number of observations		1		24		180		1,594		13,320/13,233		1,424,907/2,130,544	

[a] See footnote of Table 1

standard error. In both tables, the reported results for the standardized height and weight measures are adjusted for the natural inequality.

By construction, the proportion of between-group inequality decreases as we take more aggregated groups. In particular, the proportion of between-group to the total inequality is equal to zero at the national level (because there is only one group) and one at the level of unit records (because there is no inequality within a group).

For most of the inequality measures, the between-group component is the highest for consumption, followed by standardized height and standardized weight. At the village level, the between-group component explains a sizable portion of total inequality for consumption, but the same cannot be said about standardized height or weight. Therefore, geographic targeting is likely to be much more effective for anti-poverty programs than for child nutrition programs.

By comparing the changes in the share of between-group inequality, we find that the inequality across villages within the same commune is large compared to the inequality across communes within the same district or inequality across districts within the same province. This observation suggests that policy-makers should carefully consider the heterogeneity within a given commune whenever possible, even though commune is a reasonably disaggregated unit.

How do the numbers in Table 1 compare with a cross-country decomposition? While this study only provides results for one country and thus cannot be generalized, it does seem to indicate that health inequality is a much more local phenomenon than consumption inequality. As noted earlier, the between-country component accounts for only 31% of total world health inequality [25] as measured by $GE(0)$ for standardized height adjusted for natural inequality. In Cambodia, only 24% of the total health inequality can be explained by the between-village component of inequality using the same measure. The comparable figures for wealth inequality are much higher. More than half of the total income inequality in the world is attributable to between-country inequality. Similarly, more than half of the total consumption inequality in Cambodia is attributable to the between-village inequality.

One could object to a comparison of this sort on the grounds that the choice of reference age and sex groups—which we need to construct the standardized height and weight—is arbitrary. Furthermore, the choice of a reference group also affects the natural inequality $GE_{Natural}(\alpha)$ and the correction factor f.

However, as discussed in Fujii [18], small-area estimation of the prevalence of undernutrition is not very sensitive to the choice of the reference age and sex groups. Furthermore, as $GE_{Natural}(\alpha)$ is very small relative to $GE^{B}_{Cambodia}(\alpha)$ regardless of the reference group, the adjustment for the natural inequality does not alter the main conclusion that health inequality is predominantly a within-location phenomenon. In fact, using 6-month-old girl, 60-month-old girl, and 24-month-old boy as alternative reference groups, the maximum changes in the correction factor from the base reference population of 24-month-old girls for various values of α were 0.6 percentage points for height inequality and 5.8 percentage points for weight inequality.

We have also argued that the problem associated with the dependence on the reference group could be avoided altogether by conducting a variance decomposition of $P(0)$. As shown in Table 2, the qualitative nature of the variance decomposition results is indeed the same as the decomposition of the generalized entropy measures. The share of the between-location component of the variance in $P(0)$ is much larger for the per capita consumption measure than the standardized height and weight measures. For example, at the commune level, the proportion of the variance in the prevalence of stunting and underweight (i.e., $P(0)$ for standardized height and weight, respectively) is 6.37 percent and 5.31 percent, respectively. The corresponding figure for poverty rate (i.e., $P(0)$ for consumption) is 19.06 percent.

Using $P(0)$ for inequality decomposition has an additional advantage: As Table 2 shows, the standard errors for $P(0)$ are generally much smaller than those for $GE(0)$, especially for the consumption measure. This is because $P(0)$ is insensitive to the tails.

Thus far, we have mainly considered $GE(\cdot)$ and $P(\cdot)$ when their parameter values are equal to zero. While a parameter value of zero has some advantages over other values, as we have already argued, it is important to verify how much our results are driven by this choice of the parameter value. As we can see from Table 1, the choice of the parameter value for the generalized entropy measure does not matter very much within the range between -1 and 1. Similarly, the choice of the parameter value for the FGT measure does not affect the results a great deal for the values we considered.

7 Discussion

We have conducted decomposition analyses of health and wealth inequality by applying small-area estimation methods. The descriptive statistics provided in this study are of interest because the share of between-location inequality for small-areas was not very well known previously. Our decomposition exercise is also useful for understanding the importance of geographic information in explaining total inequality. In Cambodia, geographic targeting is potentially useful to address poverty, but its effectiveness to address child undernutrition is limited because the child nutritional outcomes are diverse even in relatively small geographic units such as villages and communes.

We have first decomposed the inequalities in health and wealth following the approach developed by Pradhan et al. [25]. By comparing the share of between-country inequality in total world inequality with the share of between-village inequality in total inequality in Cambodia, we have found that a sizable proportion of wealth inequality is determined by geography, whereas health inequality is intrinsically a local phenomenon.

We have checked the robustness of this observation in three ways. First, we checked the robustness with respect to the choice of the reference group. Second, we have also conducted a variance decomposition of the $P(0)$ index to completely avoid the arbitrary choice of a reference group and to address the comparability issue. Third, we have also checked the sensitivity of our results with respect to the choice of parameter values for $GE(\cdot)$ and $P(\cdot)$. In all of these cases, we find that our main conclusion is robust.

Because $P(0)$ at the unit-record level is a binary indicator variable, the variance decomposition of $P(0)$ allows us to compare the spatial concentration of various indicators, including mortality, literacy, unemployment, and the prevalence of diseases. Therefore, by combining this approach with small-area estimation methods, we can compare the spatial disparities across multiple dimensions of welfare and assess the potential effectiveness of geographic targeting in each of these dimensions.

References

1. Bloom, D.E., Canning, D.: The health and wealth of nations. Science **287**(5456), 1207–1208 (2000)
2. Deaton, A.: Health, inequality, and economic development. J. Econ. Lit. **41**(1), 113–159 (2003)
3. Demombynes, G., Özler, B.: Crime and local inequality in South Africa. J. Dev. Econ. **76**(2), 265–292 (2005)
4. Dibley, M.M., Gloldby, J.B., Staehling, N.W., Trowbridge, F.L.: Development of normalized curves for the international growth reference: Historical and technical considerations. Am. J. Clin. Nutr. **46**(5), 736–748 (1987a)
5. Dibley, M.M., Staehling, N.W., Nieburg, P., Trowbridge, F.L.: Interpretation of Z-score anthropometric indicators derived from the international growth reference. Am. J. Clin. Nutr. **46**(5), 749–762 (1987b)
6. Elbers, C., Fujii, T., Lanjouw, P., Özler, B., Yin, W.: Poverty alleviation through geographic targeting. J. Dev. Econ. **83**(1), 198–213 (2007)
7. Elbers, C., Lajouw, P., Mistiaen, A., Özler, B., Simler, K.: On the unequal inequality of poor communities. World Bank Econ. Rev. **18**(3), 401–421 (2004)
8. Elbers, C., Lanjouw, J.O., Lanjouw, P.: Welfare in Villages and Towns: Micro-level Estimation of Poverty and Inequality. Policy Research Working Paper 2911. The World Bank, Washington D.C., USA (2002)
9. Elbers, C., Lanjouw, J.O., Lanjouw, P.: Micro-level estimation of poverty and inequality. Econometrica **71**(1), 355–364 (2003)
10. Elbers, C., Lanjouw, P., Mistiaen, J.A., Özler, B.: Reinterpreting between-group inequality. J. Econ. Inequal. **6**, 231–245 (2008)
11. Elo, I., Preston, S.H.: Educational differentials in mortality: United States, 1979–85. Soc. Sci. Med. **42**(1), 47–57 (1996)
12. Firebaugh, G.: The trend in between-nation income inequality. Annu. Rev. Sociol. **26**, 323–339 (2000)
13. Foster, J., Greer, J., Thorbecke, E.: A class of decomposable poverty measures. Econometrica **52**(3), 761–766 (1984)
14. Fujii, T.: Commune-level estimation of poverty measure and its application in Cambodia. In: Kanbur R., Venables A.J., Wan, G. (eds.) Spatial Disparities in Human Development: Perspectives from Asia, pp. 289–314. United Nations University Press (2006)
15. Fujii, T.: To use or not to use?: poverty mapping in Cambodia. In: Bedi, T., Coudouel, A., Simler, K. (eds.) More than a Pretty Picture: Using Poverty Maps to Design Better Policies and Interventions, pp. 125–142. The World Bank (2007)
16. Fujii, T.: How well can we target aid with rapidly collected data? Empirical results for poverty mapping from Cambodia. World Dev. **36**(10), 1830–1842 (2008)
17. Fujii, T., Roland-Holst, D.W.: How does Viet Nam's accession to the world trade organization change the spatial incidence of poverty?" In: Nissanke, M., Thorbecke, E. (eds.) Globalization and the Poor in Asia: Can Shared Growth be Sustained? pp. 47–89. Palgrave Macmillan (2008)
18. Fujii, T.: Micro-level estimation of child undernutrition indicators in Cambodia micro-level estimation of child malnutrition indicators and its application in Cambodia. World Bank Econ. Rev. **24**(3), 520–553 (2010)
19. Kanbur, R.: The policy significance of inequality decompositions. J. Econ. Inequal. **4**, 367–374 (2006)
20. Kawachi, I., Kennedy, B.P., Lochner, K., Prothrow-Stith, D.: Social capital, income inequality and mortality. Am. J. Publ. Health **87**(9), 1491–1498 (1997)
21. Marmot, M., Wilkinson, R.G.: Psychosocial and material pathways in the relation between income and health: A response to Lynch et al. Br. Med. J. **322**, 1233–1236 (2001)
22. National Institute of Statistics: Cambodia Socio-Economic Survey 1997: Technical Report on Survey Design and Implementation. National Institute of Statistics, Ministry of Planning, Royal Government of Cambodia, Phnom Penh, Cambodia (1998)
23. National Institute of Statistics: General Population Census of Cambodia 1998: Village Gazetteer. National Institute of Statistics, Ministry of Planning, Royal Government of Cambodia, Phnom Penh, Cambodia (2000)
24. National Institute of Statistics: Directorate General for Health, and ORC Macro: Cambodia Demographics and Health Survey 2000. National Institute of Statistics, Ministry of Planning, Royal Government of Cambodia, Phnom Penh, Cambodia (2001)

25. Pradhan, M., Sahn, D.E., Younger, S.D.: Decomposing world health inequality. J. Health Econ. **22**, 271–293 (2003)
26. Pritchett, L., Summers, L.H.: Wealthier is healthier. J. Hum. Resour. **31**(4), 841–868 (1996)
27. Shorrocks, A.F.: The class of additively decomposable inequality measures. Econometrica **48**(3), 613–625 (1980)
28. Shorrocks, A.F.: Inequality decomposition by population subgroups. Econometrica **52**(6), 1369–1385 (1984)
29. van Doorslaer, E., Wagstaff, A., Bleichrodt, H., Calonge, S., Gerdtham, U-G., Gerfin, M., Geurts, J., Gross, L., Häkkinen, U., Leu, R.E., O'Donnell, O., Propper, C., Puffer, F., Rodríguez, M., Sundberg, G., Winkelhake, O.: Income-related inequalities in health: some international comparisons. J. Health Econ. **16**(1), 93–112 (1997)
30. Wagstaff, A.: Inequality decomposition and geographic targeting with applications to China and Vietnam. Health Econ. **14**, 649–653 (2005)
31. Wagstaff, A., van Doorslaer, E.: Income inequality and health: what does literature tell us? Annu. Rev. Public Health **21**, 543–567 (2000)
32. Waterlow, J.C., Buznina, R., Keller, W., Lane, J.M., Nichaman, M.Z., Tanner, J.M.: The presentation and use of height and weight data for comparing the nutritional status of groups of children under the age of 10 years. Bull. World Health Organ. **55**(4), 489–498 (1977)
33. WHO Working Group: Use and interpretation of anthropometric indicators of nutritional status. Bull. World Health Organ. **64**(6), 929–941 (1986)
34. WHO Working Group: An evaluation of infant growth: the use and interpretation of anthropometry in infants. Bull. World Health Organ. **73**(2), 165–174 (1995)
35. Wilkinson, R.G.: Income distribution and life expectancy. Br. Med. J. **304**, 165–8 (1992)
36. Wilkinson, R.G.: Unhealthy Societies: The Affliction of Inequality. Routledge (1996)
37. Wilkinson, R.G.: Income inequality summarizes the health burden of individual relative deprivation. Br. Med. J. **314**, 1727–1728 (1997)
38. Wright, J.A., Ashenburg, C.A., Whitaker, R.C.: Comparison of methods to categorize undernutrition in children. J. Pediatr. **124**(6), 944–946 (1994)