IARIW 2025

IARIW 2025

Thursday, October 2 & Friday, October 3

Asset-wealth and malnutrition estimations in Lao PDR: The role of remote sensed data

Diana Carolina Garcia Rojas

(World Bank)

Paper prepared for the IARIW–World Bank–UEB/VNU Conference on “Improving Well-being Measurement in Data-challenged Environments in Developing Countries for Better Evidence-based Policies” October 2-3, 2025

Session 2B Data and Methods
Time Slot: Thursday, October 2, 1:15- 3:15 PM

Asset-wealth and malnutrition estimations in Lao PDR: The role of remote sensed data

Author: Diana C. Garcia Rojas

Abstract

In low- and middle-income countries, monitoring poverty and chronic malnutrition at a more granular geographical level than the national or regional level presents a significant challenge. Data capturing households' socioeconomic conditions is often scarce and difficult to access due to concerns about anonymity. Academics and policymakers have employed statistical techniques, such as small area estimation (SAE), to estimate poverty and malnutrition at more disaggregated levels. However, these techniques require large-scale population data, such as census data, which is typically collected only once every ten years in most countries. Recently, broader access to satellite imagery has opened new opportunities to estimate household welfare in small areas during off-census years. In this paper, I test the use of area-level and unit-level models to estimate welfare indicators using only aggregated data in a rural-dominant country, such as Lao PDR. These estimations consider the exclusive use of aggregated remotely sensed data and the potential gains from incorporating census aggregates. I then apply the best-performing models to estimate average asset wealth and the prevalence of stunting at the district level in 2016. The main findings are as follows: First, estimations using only remotely sensed data are generally more suitable for urban areas, but they still do not outperform those that incorporate census aggregates. Second, a unit-level model using clusters as units can perform as well as an area-level model when the outcome is linear, such as average asset wealth, but the former offers a more balanced trade-off between variance and bias than the latter. Finally, I find that using estimations of average asset wealth as a proxy to estimate stunting rates can improve malnutrition estimations, which are generally less precise, but are not sufficient.

Introduction

Socioeconomic indicators, like the Sustainable Development Goals (SDGs) 1 and 2, which are related to poverty and malnutrition, are hard to monitor at highly disaggregated geographical levels mainly due to data availability or anonymity concerns. Estimations of households' conditions are usually only available at the national level or in large geographical areas like provinces or regions due to the high cost of carrying out surveys or censuses. Surveys are generally collected sporadically and do not have enough statistical power to estimate welfare indicators in small areas, like villages or districts. In the past few decades, statistical techniques, like small area estimation, were developed to estimate poverty and malnutrition at these lower geographical levels (Battese et al., 1988; Elbers et al., 2003; Molina et al., 2022).

Small area estimation (SAE) techniques have become essential for producing reliable poverty and malnutrition estimates in data-scarce environments. Area-level models, such as the Fay-Herriot model, are widely used due to their ability to incorporate aggregated auxiliary data and improve the precision of direct survey estimates. In contrast, unit-level models offer an alternative that can reduce variance by leveraging more granular data and can be more precise at smaller areas than area-level models. Molina et al (2014) demonstrated through simulation and real-world applications that empirical best prediction (EBP) estimators, which are based on unit-level models, provide more efficient estimates than direct estimation, especially for continuous outcomes. Further supporting this, Jiang and Lahiri (2006) highlighted that unit-level models allow for greater flexibility in incorporating complex covariate structures, which translates into improved predictive accuracy when auxiliary data is granular and reliable. These findings underscore the advantage of unit-level approaches in contexts where detailed covariates are available, while also emphasizing the importance of accounting for survey design features to maintain unbiasedness. However, these models rely on the existence of full population data, like censuses, which are collected only every 10 years or so in low and middle-income countries.

Nowadays, due to the higher need for timely and more disaggregated socioeconomic data, the focus has shifted to the use of remote sensed data and satellite imagery to estimate welfare indicators. A new trend in literature has attempted to estimate levels of poverty or wealth in census and off-census years for small areas, even at pixel level, using sophisticated statistical or machine learning techniques (Chi et al., 2022; Jean et al., 2016; Khachiyan et al., 2022; Lee & Braithwaite, 2022; Yeh et al., 2020). Several studies using SAE show that incorporating geospatial covariates such as night-time lights, vegetation indices, and land cover improves model precision compared to survey-only approaches (Newhouse, 2024). For example, Permatasari and Ubaidillah (2025) demonstrated that Fay-Herriot models using remote sensing data yielded more accurate poverty estimates for unsampled areas than direct estimates. Nonetheless, various critiques arise relate to the non-rigorous, or lack of, validation of these methods. Most of these studies use the same survey data to validate the techniques and rely on proxies that may not fully capture welfare dynamics (Corral et al., 2022, 2025; Van Der Weide et al., 2024). Also, using machine learning tools like deep learning makes it hard to replicate the estimations due to the complexity of the technique (Machicao et al., 2022).

This paper aims to contribute to the literature by: (1) validating the use of area-level and unit-level models to estimate welfare indicators when household data are unavailable and only remotely sensed or aggregated administrative data are available in rural settings such as Lao PDR; and (2) applying the best available methods and data to estimate asset wealth and stunting rates at the smallest possible geographic level for 2011 in Lao PDR. The motivations for pursuing this study in a country like Lao PDR are twofold: first, it serves as a representative example of an agriculture-dependent country with low urban density; second, it offers suitable data to carry out the analysis and apply the best practices.

Lao PDR has had three DHS surveys rounds, for the years of 2011, 2017 and 2023, for which the clusters geolocation is known for the first two. Also, in 2011, the government increase the number of land concessions given for agricultural production in certain areas(Asian Development Bank, 2017). Therefore, precise estimations of welfare indicators at district or village level, calculated from DHS surveys, could be ideal to track the effect of these policies. Garcia Rojas et al., (2024) used a unit-level model to estimate stunting rates at district level for 2015 using the population census from the same year and found that average asset wealth was one of the most important predictors in the model. Both of these indicators could be estimated for 2011, however, there is no population census available close to this year, Thus, a SAE area-level or unit-level model that relies only on aggregated remote sensing data will be a potential solution to estimate these indicators at the district level. Therefore, this paper attempts to validate which method and data are more suitable to have precise welfare estimates in a rural country like Lao PDR using surveys like the DHS.

The methodology of this study consisted of three steps:

1. Validation of the method and data to be used: A synthetic asset-wealth vector was constructed using census data, and 100 synthetic surveys were sampled using a similar methodology to the DHS. This synthetic data was used to test whether an area-level model (the Fay-Herriot model) is better suited than a unit-level model (Pseudo EBP model) that uses cluster as units to estimate average asset-wealth and

poverty rate[1]. Also, it tested the efficiency of using remote sensed data compared to census aggregated covariates in the case of Lao PDR.

2. <u>Estimation of asset-wealth poverty at the smallest area</u>: Using the best practices identified in the first step, asset-wealth poverty rates were estimated for the 2017 survey.
3. <u>Estimation of stunting at the smallest area</u>: Considering asset-wealth is one of the best predictors of stunting, stunting rates were estimated using the estimations of average asset-wealth as proxy.

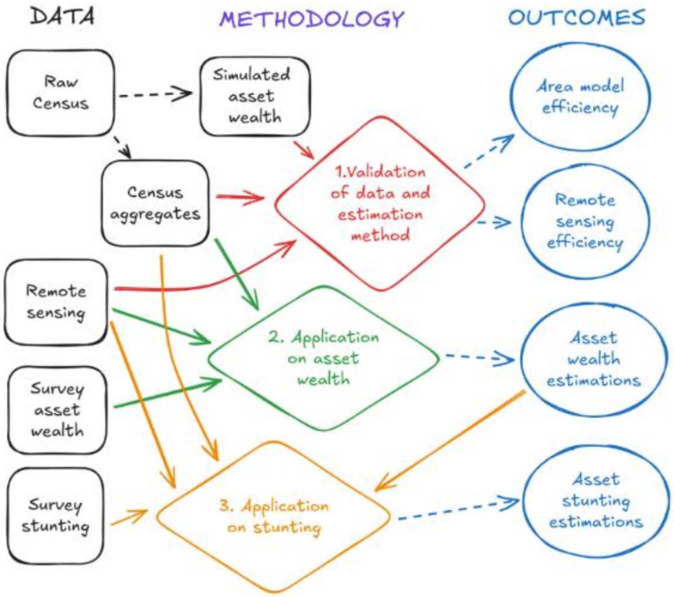

1. Summary of the methodology

The key findings are: first, models that rely solely on remotely sensed data tend to perform better in urban settings. However, they still fall short compared to models that incorporate census-based aggregates. Second, when estimating linear outcomes like average asset wealth, unit-level models that use clusters as the unit of analysis can match the performance of area-level models. Moreover, unit-level models offer a more favorable balance between bias and variance, as the source of potential errors. And lastly, using estimates of average asset wealth as a proxy for stunting rates can enhance the accuracy of malnutrition estimates, which are typically less reliable on their own, but are not sufficient to increase the accuracy of the estimates.

---

[1] This non-monetary poverty rate was defined by the percentage of households that have a lower asset wealth index that a specific threshold in the standardized wealth composite. This threshold was set to match the national monetary poverty rate estimated by the Lao Expenditure and Consumption Survey (LECS) from 2018.

Data

Population data

The two main data sources for this study are the Population and Housing Census of 2015 and Lao PDR Demographic and Health Survey (DHS) of 2017, or more specifically the Lao Social Indicator Survey (LSIS II). The population census was used to create the synthetic asset wealth index and construct the census aggregates at village or district level, as well as for the imputation of the area-level and unit-level models. It includes features for all households and individuals in the country, such as the average household size, dependency ratio, literacy rate, main dwelling material (roof, floor, walls), main water source, and the percentage of the population under/over certain ages, among others. The DHS, which collects data on asset wealth and the prevalence of stunting, was used to apply the estimation methods. The objective of this survey is to provide information on social and economic household conditions, with a strong focus on women's health and child nutrition. DHS surveys are generally open-source, and in some countries, clusters are geo-referenced, which is the case for Lao PDR in 2011 and 2017.

GIS data

This study used not only the two main sources described above but also a wide range of remote sensing and satellite data for Lao PDR. The original aggregated remote sensing data came from SERVIR MEKONG, NASA/NOAA, OpenStreetMap (OSM), and MOSAIKs. Given that more than 50% of the labor force in Lao PDR works in the agricultural sector, it was important to incorporate processed remote sensing data related to land cover and land use. For example, the land cover dataset developed by SERVIR MEKONG includes raster information on land types such as cropland and orchard plantations. Additionally, spatial data capturing shifting cultivation, such as that developed by Chen et al., (2023), can improve model predictability compared to data based solely on temperature or precipitation, as used in previous studies. Nonetheless, commonly used spatial features such as night-time lights, urban built-up areas, elevation, and precipitation were also included. Furthermore, OpenStreetMap data allowed the identification of facilities (e.g., schools, hospitals, banks, places of worship) within a target area, such as a village or district, as well as the distance from urban hubs to the nearest amenities, such as markets, for the study year (2017). Although MOSAIKs features were considered in this study, they contributed less to model performance compared to the other sources described above.

Methodology

Synthetic asset wealth index

The procedure for building the synthetic wealth index follows the steps outlined in the LSIS and DHS documentation (Lao Statistical Bureau, 2017) and uses the 2015 Lao PDR

Population Census as the primary source. Although not all the information collected by DHS to construct the wealth index is available in the census, the synthetic wealth index aims to measure households' relative economic status using household assets and living conditions, similar to the DHS approach.

The DHS wealth index is a composite measure constructed using principal components analysis, where factor scores are derived from standardized data on asset ownership, housing characteristics, access to water, sanitation, and other variables. To minimize urban bias, factor analysis is performed on the overall sample as well as separately for urban and rural households. The final combined factor is then a linear combination of the urban or rural factor and the full-sample factor[2]. All households are assigned a wealth score based on their assets and characteristics, using these combined factor scores

The household data used specifically for the LSIS II wealth index is the following:

> "(...) Main material of dwelling floor, roof and external walls; possession by the household of a fixed telephone line, a radio, a clock , a sofa/wooden settee, a bed/mattress , electricity, a television, a refrigerator a fan, a water pump , an air-conditioner, a washing machine, a CD/DVD player/home theater, an iron, a rice cooker/steamed cooker, a watch , a bicycle, a motorcycle or scooter, an animal-drawn cart , a car, truck or van, a boat with a motor, a tak tak , a computer or a tablet, a mobile phone, internet at home , agriculture land, livestock, herds other farm animals or poultry, a bank account ; type of the cookstove, type of fuel or energy source used for the cookstove and location where the cooking is done; space heating, type of fuel and energy used for space heating; what is used to light the household; source of drinking water; location of water source; reasons for insufficient quantity of water ; type and location of sanitation facility, sharing of sanitation facilities ; place for handwashing and availability of soap."

Unfortunately, the sources underlined in the previous paragraph were not available in the 2015 Census. However, this study included other variables that are not listed in the LSIS II description but are commonly included in other DHS calculations: tenure status of the household, area occupied, and number of rooms (Rutstein & Johnson, 2004).

---

[2] The full, urban and rural factors were validated by comparing the principal component with ownership of certain goods, like television or motor vehicles. In the case of the full and urban factors, the vector had to be rotated (multiplied by -1) in order to coincide with the concept of welfare.

Sampling procedure

Once every household in the census was assigned a wealth score, 100 synthetic surveys were sampled from this population. The sample design follows the specifications outlined in Appendix A of the LSIS II documentation. The full sample size comprises 23,400 households,[3] and the minimum sample size for each province (explicit strata) was 1,100 households. The sample was allocated to each province in proportion to the square root of the number of households in the area and then adjusted to meet the minimum requirement. The number of households per cluster (village) was set at 20 resulting in the selection of 1,170 villages overall[4].

Using systematic probability proportional to size (PPS) and systematic sampling, villages were sampled from a list ordered by village category to obtain proportional allocation over the implicit strata (urban, rural with road and rural without road). An overview of the sample allocation used in this study is presented in Appendix XX, and it shows similar numbers as those detailed in the LSIS II final report[5]. A difference between the procedure followed by this study and the LSIS was how the selection of households was done within the cluster. The random systematic selection is performed over all the households in the census and not the listings gathered by the LSIS survey team for each province, which are not openly available. Also, segmentation of clusters was not performed and there was no adjustments for non-response, therefore the calculation of weights simply considers the probability of selection (p_{sc}) in two stages:

$$W_{sc} = 1/(p_{1sc} * \ p_{2sc})$$

The first stage refers to the selection of villages/clusters (c) within the province/strata (s) and the second stage to the selection of households within the cluster. The probabilities are then calculated as follows:

$$p_{1sc} = \frac{n_s * M_{sc}}{M_s} \ \ (or\ 1\ if\ selected\ with\ certainty)$$

$$p_{2sc} = \frac{20}{M_{sc}}$$

Where n_s is the number of clusters sampled in the province and M_{sc} and M_s are the total number of households in the cluster and the province respectively. These probabilities will

---

[3] The overall sample size was decided by considering the statistical power needed to measure underweight prevalence among children age 0-4 years in the LSIS II.

[4] This number was selected based on design effect, budget, and personnel constraints.

[5] The R code used to performed the sampling process is also available at request.

be relevant for the PEBP model specifications, which uses villages as units, explained in the next subsections.

Area-level model at village and at district level

The first small area estimation methodology considered in this study is the most used area-level models, best known as Fay-Herriot models, or FH from now (Fay & Herriot, 1979). These models are designed to improve direct survey estimates for areas or domains where they usually fail due to small sample sizes. The method borrows strength from auxiliary data at the desired area level and relies on the following linear mixed model:

$$\widehat{\theta_i} = x_i * \beta + u_i + e_i$$

Where $\widehat{\theta_i}$ is the direct survey estimate, x_i is the auxiliary data, u_i are area-specific random effect and e_i the sampling error. Usually, the parameters β are estimated using empirical best linear unbiased prediction (EBLUP) and the resulting FH estimator is the weighted average of the direct estimator and the model prediction:

$$\theta_i^{\widehat{FH}} = \widehat{\lambda_i} * \theta_i^{\widehat{Dir}} + \left(1 - \widehat{\lambda_i}\right) * x_i * \hat{\beta}$$

This estimator uses a shrinkage factor $\widehat{\lambda_i} = \frac{\widehat{\sigma_u^2}}{\widehat{\sigma_u^2} + \widehat{\sigma_e^2}}$ that is larger when the sampling variance is small, giving more weight to the direct estimator. The model assumes sampling variances are known, but this is rarely the case, and they are commonly computed using survey design-based variance. In this study, the sampling variances were calculated using the R package survey when the indicator was a simple average or by a built function when the indicator was a ratio. This built function used Jeffrey's smoothing and arcsin transformation.

The FH model was tested to estimate several indicators at village and district level using the synthetic survey samples from the constructed wealth index described above. All covariates were aggregated to the specific level of estimation and the model covariates were selected in three steps. First, variables that did not match distributions in survey and population using chi-square tests were dropped. Second, variables with larger coefficients after using elastic net were selected. Third, variables were dropped if they incur in collinearity. Finally, the FH model was calculated using the function fh from the emdi package and each specification is described in the results section.

Unit-level model using cluster estimates as units

The objective of this study is to find the best method to estimate at small areas when individual or household data is missing and only aggregated data at a geographical level is available. FH models tend to be unbiased, because they use direct survey estimates, but the variance of the estimator can be large due to only relying on aggregated auxiliary data at the desired level of estimation. A unit-level model is then considered to be another useful SAE model in this set up. Due to its potential to borrow strength from a more detailed dataset, unit-level models usually produce estimates with lower variance than FH. The formal specification of a unit-model is the following:

$$y_{ij} = x_{ij}^{T} * \beta + u_i + \varepsilon_{ij}$$

Where y_{ij} is the indicator for unit i in area j, x_{ij} are the covariates and $u_i$and ε_{ij} are the area random effect and the unit error respectively. Common unit-level models are designed to use household or individual level data, but in this case, I attempted to use unit-level models, more specifically the pseudo empirical best prediction model (Guadarrama et al., 2018) or PEBP from now on, with aggregated data at the unit-level. The PEBP model is an extension of the Empirical Best Predictor model (Molina & Rao, 2010) or EBP, and it can be described as:

$$If \quad \bar{y}_{jw} = \frac{\sum_{i\in s_j} w_{ij} y_{ij}}{\sum_{i\in s_j} w_{ij}} \quad and \quad \bar{X}_{jw} = \frac{\sum_{j\in s_i} w_{ij} x_{ij}}{\sum_{j\in s_i} w_{ij}}$$

$$then \quad \hat{\theta}_i^{PEBP} = \bar{X}_i^{\top} \hat{\beta}_p + \hat{\gamma}_i^{(p)} \left(\bar{y}_{iw} - \bar{X}_{iw}^{\top} \hat{\beta}_p\right)$$

$$where \quad \hat{\gamma}_i^{(p)} = \frac{\hat{\sigma}_{u,p}^2}{\hat{\sigma}_{u,p}^2 + \hat{\sigma}_{\varepsilon,p}^2 / \tilde{n}_j} \quad , \quad \tilde{n}_j = \frac{\left(\sum_{i\in s_j} w_{ij}\right)^2}{\sum_{i\in s_j} {w^2}_{ij}}$$

PEBP is more suitable for the estimations in this study because it accounts for weights w_{ij} in the sample and the effective sample size $\tilde{n}_j$ for area j. Therefore, the proposal is to use the first stage of the sampling frame, or the probability of selection of villages p_{1si}, to construct the sample weights for the units in the PEBP: $w_{ij} = 1/p_{1si}$. Also, the total number of households per village will be used as population weight.

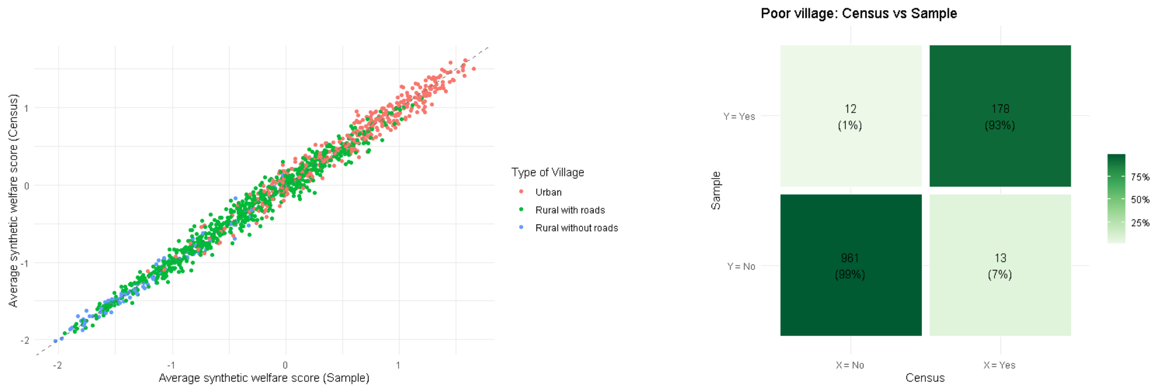

2. Weighted average of wealth score per village (left) and percentage of poor villages (right) compared from full census and a random sample.

It is important to point out that the unit is calculated as the average of the indicator at the cluster/village level, thus the unit error has measurement error by design from the second sampling stage: $\varepsilon_{ij} = e_{ij} + \partial_{ij}$. However, as presented in Figure 2, when considering the main indicators (average asset wealth index and poor village status) from our synthetic sample, the belief is that this type of error is relatively small, and this is one of the motivations to try the PEBP model in this set up. The expectation is that the PEBP estimates will have lower variance than FH but have larger bias than a formal unit-level model, where household or individual data is available. For future research, the sampling variance within the units could be accounted for by including the sampling variance in the calculation of the unit error variance as follows: $\hat{\sigma}^2_{\varepsilon,p} = \hat{\sigma}^2_{e,p} + \hat{\sigma}^2_{\partial,p}$. However, as of this date, the emdi package does not include this specification, so the results presented in the next section won't incorporate this calculation. This study chooses to only rely on existing functions from the statistical package for easy replicability of the results but acknowledges that estimations could be improved if measurement error is considered in the calculations.

Similarly to the FH model, the PEBP model was tested to estimate several indicators using the synthetic survey samples, however, in this case the estimation was only attempted at district level. All covariates were aggregated to the cluster level and the model covariates were selected in the same three steps as in FH. Finally, the PEBP model was calculated using the ebp function from the emdi package and, even though each model specification is described in the next section, all estimations relied on REML for consistency.

Results

Wealth composite and sampling procedure

The validation of the methods presented above depends on how the synthetic wealth vector designed as explained above can be a good proxy of the DHS wealth index. In Figure 3, the histogram shows how the final composite represents the distribution of households as expected if the assumption that households in rural villages without roads have lower wealth than households in urban villages. Also, we can see a clear linear relationship between the village average wealth index from the synthetic composite and the DHS wealth index.

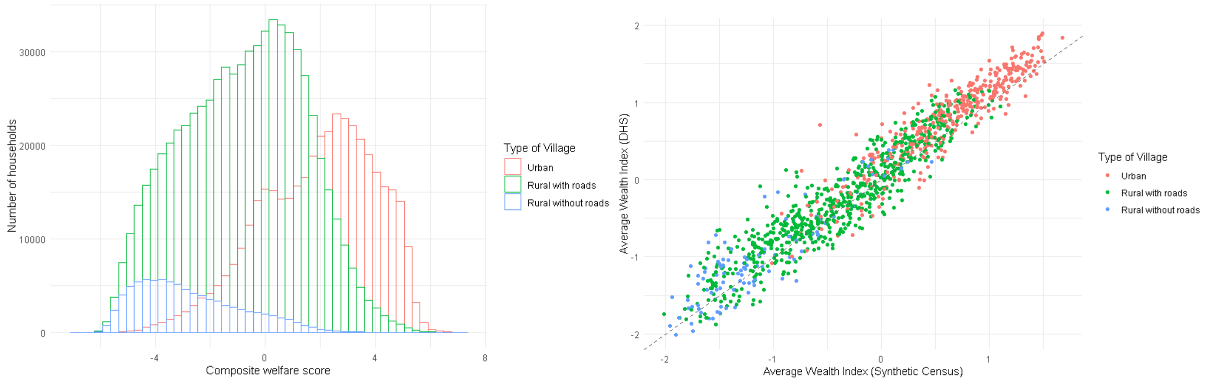

3. Distribution of synthetic wealth index across all households (left) and comparison of average village wealth index from the synthetic composite and the DHS wealth index (right).

The verification of the sampling exercise is also important if the aim is to mimic with the synthetic samples how the wealth index, as well as other indicators, are estimated in the LSIS II. Table 1 presents the sample frame by village type for the DHS survey compared to a random synthetic sample using the sampling procedure described above. The percentage of villages sampled in each survey matched exactly.

Table 1. Sample Frame used for synthetic samples compared to DHS

Village Type	**DHS (number)**	**DHS (percentage)**	**Synthetic Sample (number)**	**Synthetic Sample (percentage)**
Urban	371	32%	371	32%
Rural with roads	681	59%	691	59%
Rural without roads	106	9%	102	9%

Also, the sampling procedure followed in this study used province as the explicit strata and it is expected that the main indicators to be estimated in this study should match the distribution at this geographical level. Figure 4 shows how the survey estimates from a random synthetic survey match the true values aggregated at the province level. It is also clear that the variance is lower when calculating the average wealth score than the poverty rates, which are calculated using as poverty threshold (-1.04) that equals the quantile within the distribution that represents the national poverty rate (18.3%) according to the latest consumption survey in Lao PDR (Lao Statistics Bureau, 2023) .

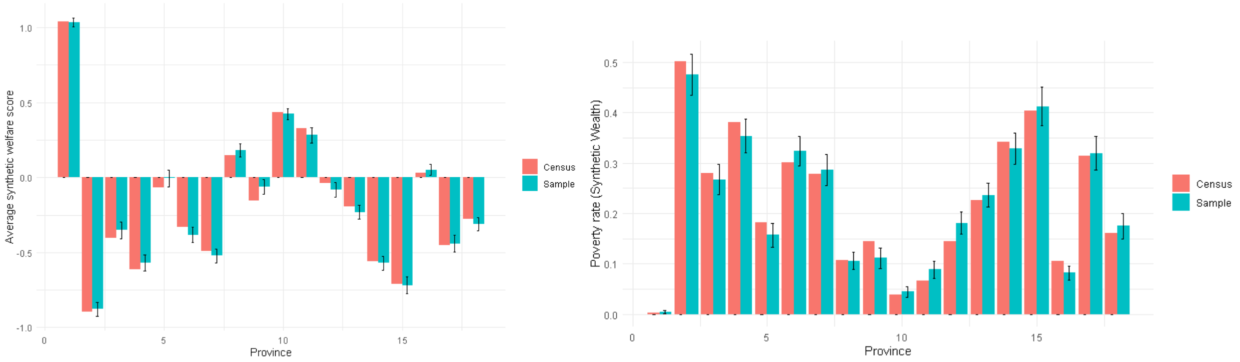

4. Average wealth score (right) and poverty rates at province level from the synthetic sample compared to the census.

Area-level model at village level

The main hypothesis to be tested in this paper is that better suitable data at a highly disaggregated level could improve the estimation of certain indicators using small area estimation. As discussed in the introduction, stunting rates have shown to be highly correlated with the asset wealth index in the LSIS II, however, in off-census years this variable would be unavailable. Therefore, a proxy of this indicator could be estimated using village or district level aggregates, like GIS or administrative data.

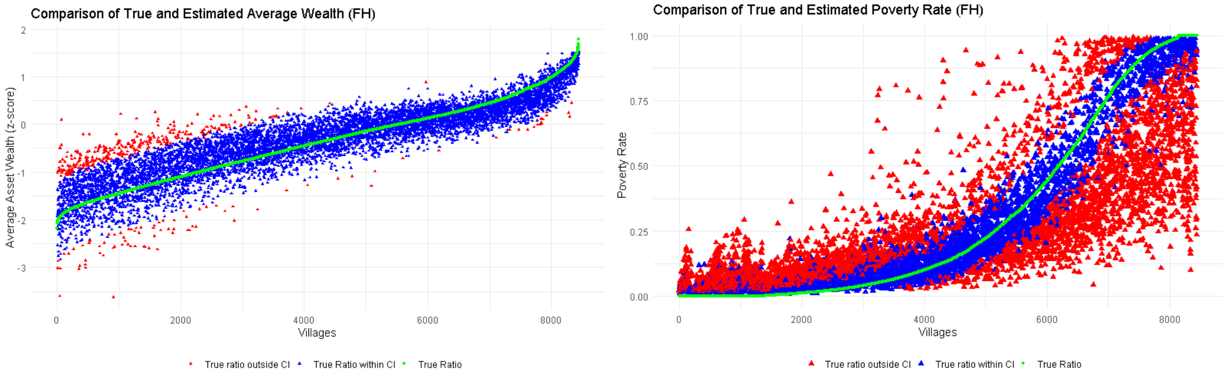
Comparison of True and Estimated Average Wealth (FH)

Comparison of True and Estimated Poverty Rate (FH)

5. Area-level village estimates of average wealth (left) and poverty rates (right)

First, using the 100 synthetic samples, I tested how the estimates from a FH model at village level compared to the true average wealth and poverty rates. The FH model used for the estimation of average wealth used no transformation of the dependent variable since the wealth composite is normal by design, but for the poverty rates[6] an arcsin transformation was used as well as the back-transformation with bias correction (Jiang et al., 2002; Sugasawa & Kubokawa, 2017). Figure 5 presents how the average estimates across the 100 surveys vary widely, especially for villages with low average wealth score and large poverty

---

[6] Poverty rates were also smooth using additive or Laplacian smoothing (Jeffrey's prior) due to the large number of villages with zero poverty rate (Sakai, 2025)

rates. This result shows how estimations at the smallest area using area-level models are highly imprecise and should not be used for policy recommendations.

This exercise also provides some evidence that relying only on GIS as auxiliary data can worsen the estimates for poorer areas. In Figure 6, the comparison between the average wealth estimates using all aggregated data (left) versus using only GIS auxiliary data (right) demonstrate that, especially in rural, and poorer areas the predictions suggest higher wealth scores than what they truly are.

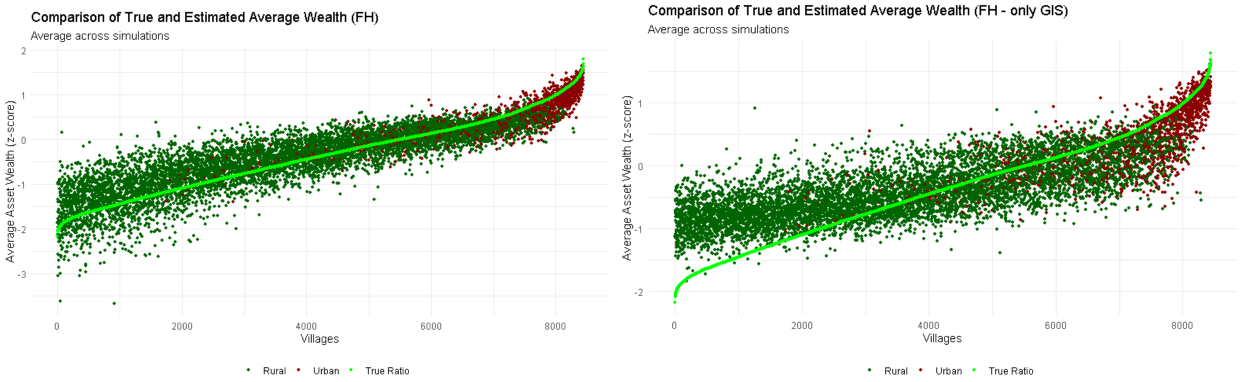

6. Area-level village estimates of average wealth using all data available (left) and only GIS data.

The premise is that GIS data would be able to capture attributes better associated to urban than rural settings. Even though this study tried to use as much open-source satellite data related to the agricultural sector, or the existence of amenities like hospitals or markets in remote areas, these features were not enough to improve the estimates for rural villages.

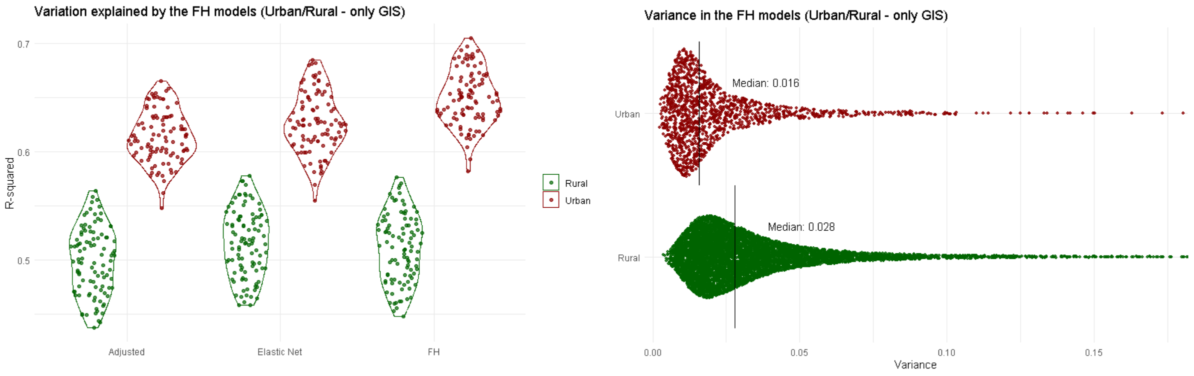

7. Variation from FH urban and rural models estimating average wealth

The results regarding explanatory power and variance from two separate FH models for rural and urban villages indicate that the GIS variables can explain better the variation of the estimates for urban villages and thus the variance of the estimates are much lower (Figure 7). Therefore, any researcher making use of only GIS variables in combination with area-level models should be aware of potential biases toward urban settings. For this study, due to the large MSE in both settings, these estimates at village level were not used as proxy to estimate stunting rates at district level. The next best alternative would be the average wealth estimates at district level.

Area-level model at district level

The advantages of using an area-level model, like FH, in this study to estimate the main indicators at a higher geographical level than the village, is that: first, it only needs district level auxiliary data and that by design, the estimates are theoretically unbiased, because they depend on the direct estimates. However, the FH model assumes that the sampling variance is known, which is usually not the case. Thus, an estimate is used using the specification of the sampling procedure and this approximation can lead to a large variance in the estimates and thus create overfitting issues.

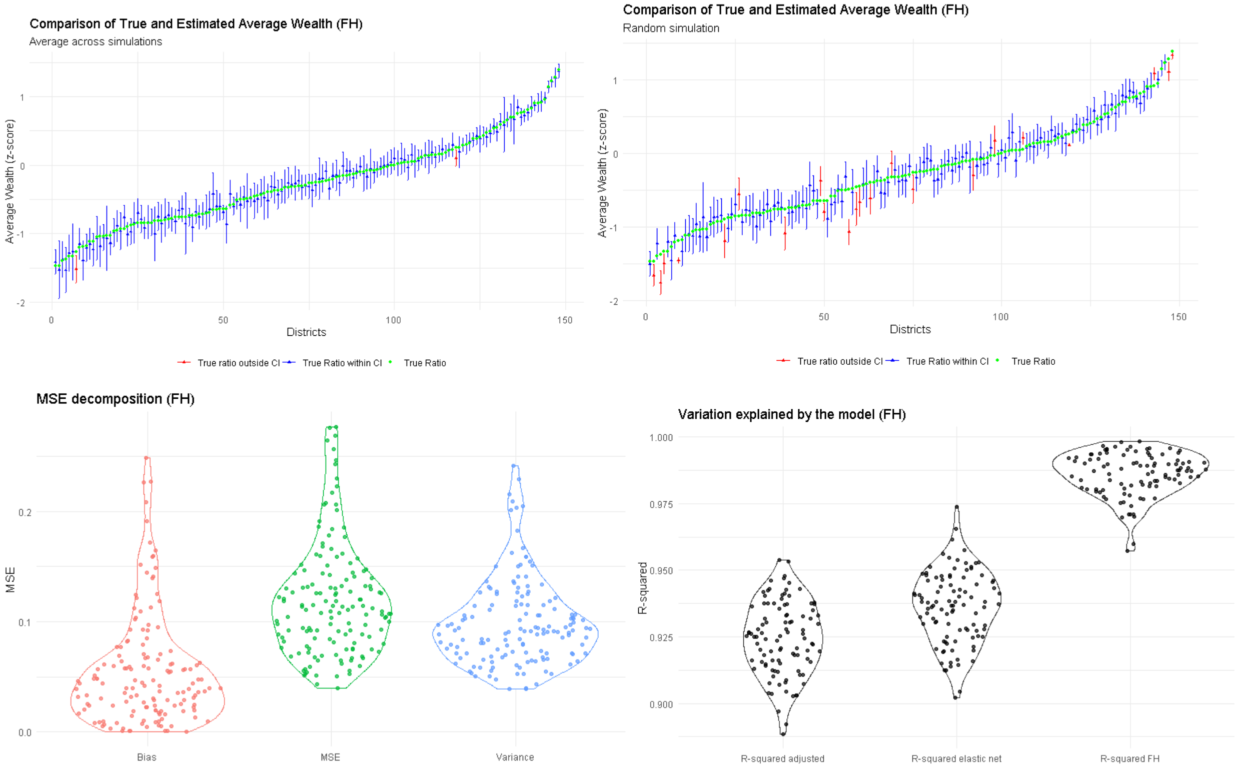

8. Validation of FH model to estimate average wealth at district level. Point estimates and confidence intervals compared to the true census aggregates are presented in the top figures for the average across samples (left) and for a random survey (right). Also, the decomposition of the MSE (bottom-left) and the R-squared (bottom-right) for this model.

The FH model used for the estimation of average wealth at district level used the simplest specification: no transformation of the outcome variable and the REML method of estimation was selected. The results presented in Figure 8 show how the area-level estimates did gain from lower bias, due also to the good explanatory power of the auxiliary data. Although, the estimates appear to be accurate when averaging across surveys, the variance is still large enough so that the estimates from a random survey do not match the true values for several districts when considering confidence intervals.

These previous findings are even more visible when the outcome is nonlinear, like the poverty ratio. The specification for this FH model used arcsin transformation and the back-transformation with bias correction as in the village level model. Figure 9 shows how the MSE is mostly affected by high variance, leading to larger errors in the estimation for a random survey. The overfitting is also strengthened by the large r-squared values. Thus, in the next subsection, I tested another method that could reduce the large variance effect by taking advantage of more granular data, at village level. A unit-level model like the PEBP can be an alternative to the FH model.

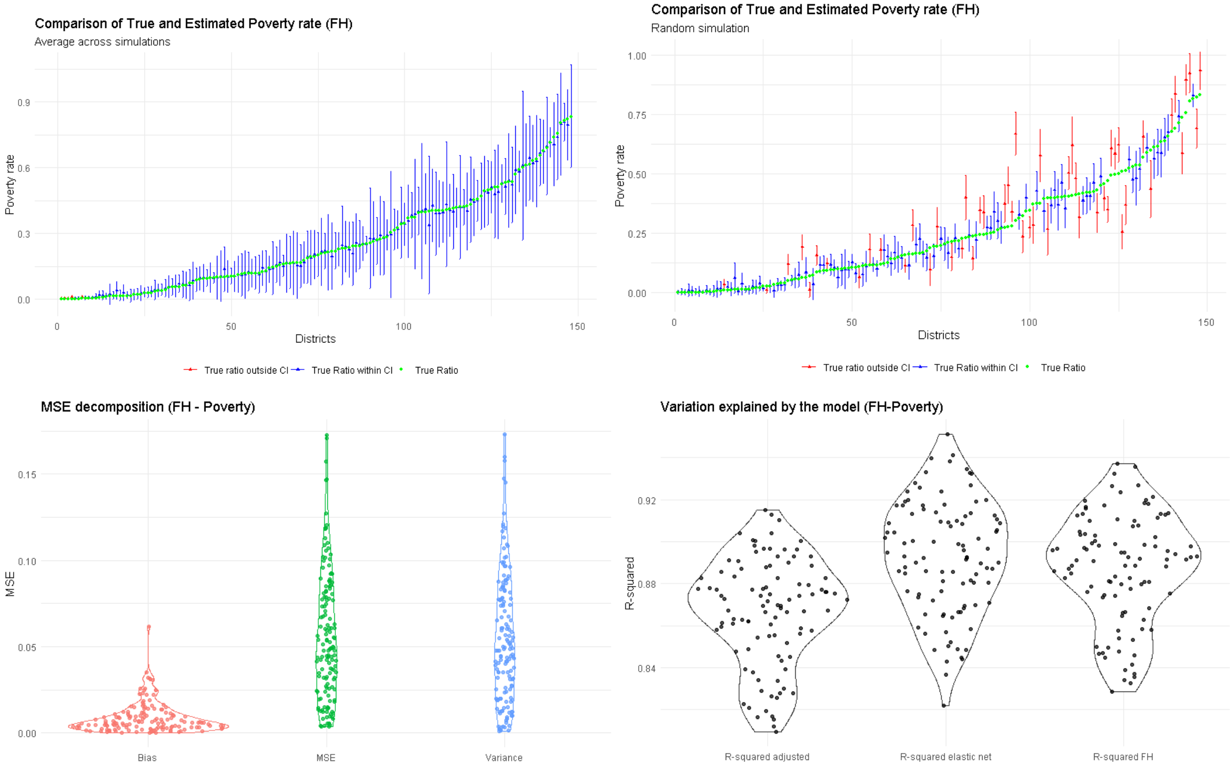

9. Validation of FH model to estimate poverty rates at district level. Point estimates and confidence intervals compared to the true census aggregates are presented in the top figures for the average across samples (left) and for a random survey (right). Also, the decomposition of the MSE (bottom-left) and the R-squared (bottom-right) for this model.

Unit-level model at district level (clusters as units)

The use of a unit-level model, with the aggregated outcome at cluster level as the unit, is only considered in this set up because of the stratified sampling used by surveys like DHS. This sampling procedure offers the possibility to use independent weights for the cluster and the households within the survey. Also, more detailed auxiliary data at village level can reduce the high variance per district found with area-level models. However, the estimation could have larger errors due to the added measurement error in the main indicator. This study validates this method using the 100 surveys sampled from the census with the synthetic wealth index vector.

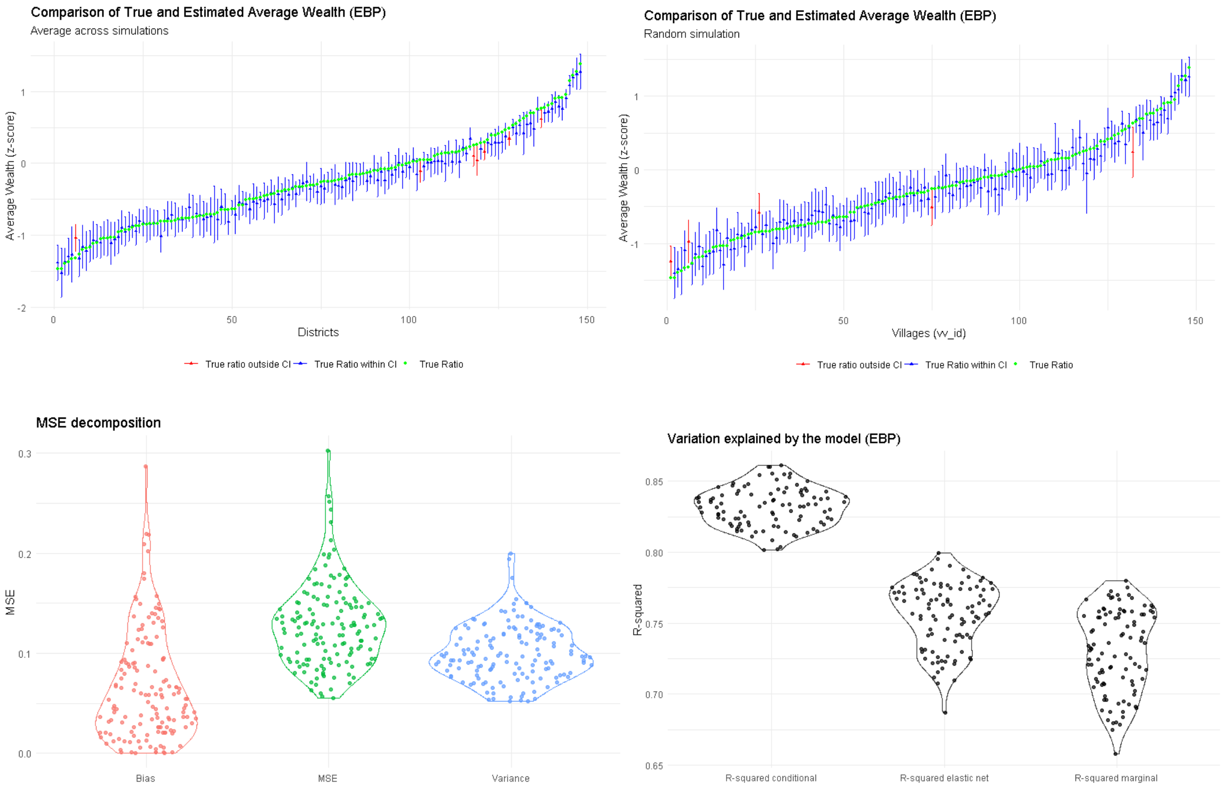

10. Validation of PEBP model to estimate average wealth at district level (using clusters as units). Point estimates and confidence intervals compared to the true census aggregates are presented in the top figures for the average across samples (left) and for a random survey (right). Also, the decomposition of the MSE (bottom-left) and the R-squared (bottom-right) for this model.

The results in Figure 10 show the decomposition of the MSE presents an expected reduction in variance, compared to the FH model, and an expected increase in bias. The PEBP model specification used for average wealth at district level used no transformation of the outcome variable and the weights were applied as described in the methodology. Even though the covariates utilized for this model do not explain the variation of the indicator as well as in the FH model, the linear models across surveys were still strong with R-squared above 65%. The estimation obtained by this unique methodology led to less overfitting when used over a random sample. However, the confidence intervals are larger than those with the FH model. Taking into consideration that on average the MSE for both models are around the same level (FH=0.17 and PEBP= 0.19), PEBP appears more suitable for constructing the average wealth proxy needed to estimate stunting, as it provides a better balance between variance and bias compared to FH.

Regarding the estimation of poverty rates (Figure 11), the PEBP also showed an improvement on overfitting, compared to the FH model, but the large bias from a potential model misspecification and a compound effect of the measurement error[7]. The latter is a common characteristic when estimating a nonlinear indicator and it is less pronounced when estimating averages like the one presented above. This method will be used to estimate stunting rates in the application subsection of the results. Nonetheless, the estimates are compared to the FH estimates to have a balance overlook.

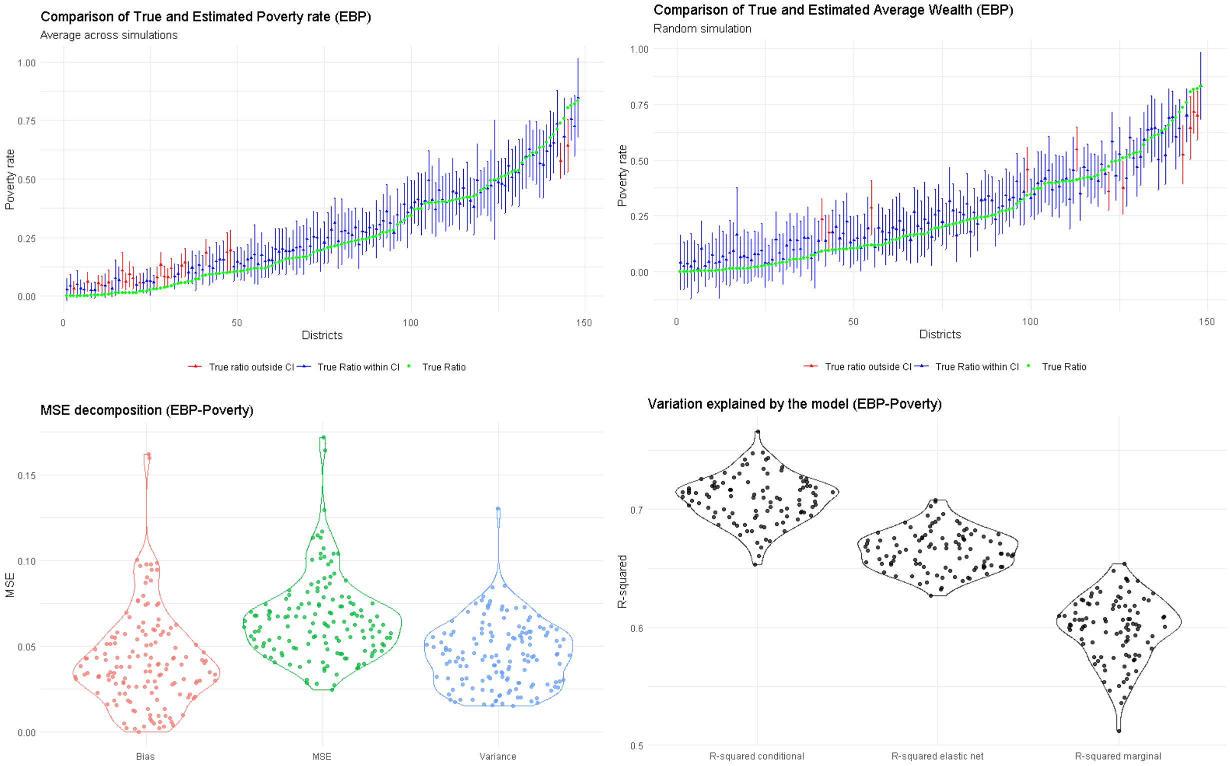

11. Validation of PEBP model to estimate poverty rates at district level (using clusters as units). Point estimates and confidence intervals compared to the true census aggregates are presented in the top figures for the average across samples (left) and for a random survey (right). Also, the decomposition of the MSE (bottom-left) and the R-squared (bottom-right) for this model.

Application: Estimation of average wealth

In the previous simulation, the method that was found to be more suitable to estimate average asset wealth index was the PEBP model at the district level aggregation. Figure 12 shows the results by first, comparing the PEBP estimates with the direct estimates using MSE and the coefficient of variation (CV) and second, comparing the estimation using all available auxiliary data and using only GIS data. The main finding from this application is that

---

[7] The specification used for this model was similar to the specification used for average wealth but the main indicator, the poverty rates at village level, was smoothen using additive or Laplacian smoothing to avoid large zero values (non-poor villages).

the estimations have more stable and accurate estimates than the direct estimates and, thus, could be used as a better proxy to estimate stunting rates. Statistical offices set thresholds for the reliability of the estimates between 16% and 33%. The average asset wealth could be reliable even when the estimation uses only GIS. This could indicate that average asset wealth could be estimated in off-census years, or in years where administrative data is not available, using PEBP models with clusters as units.

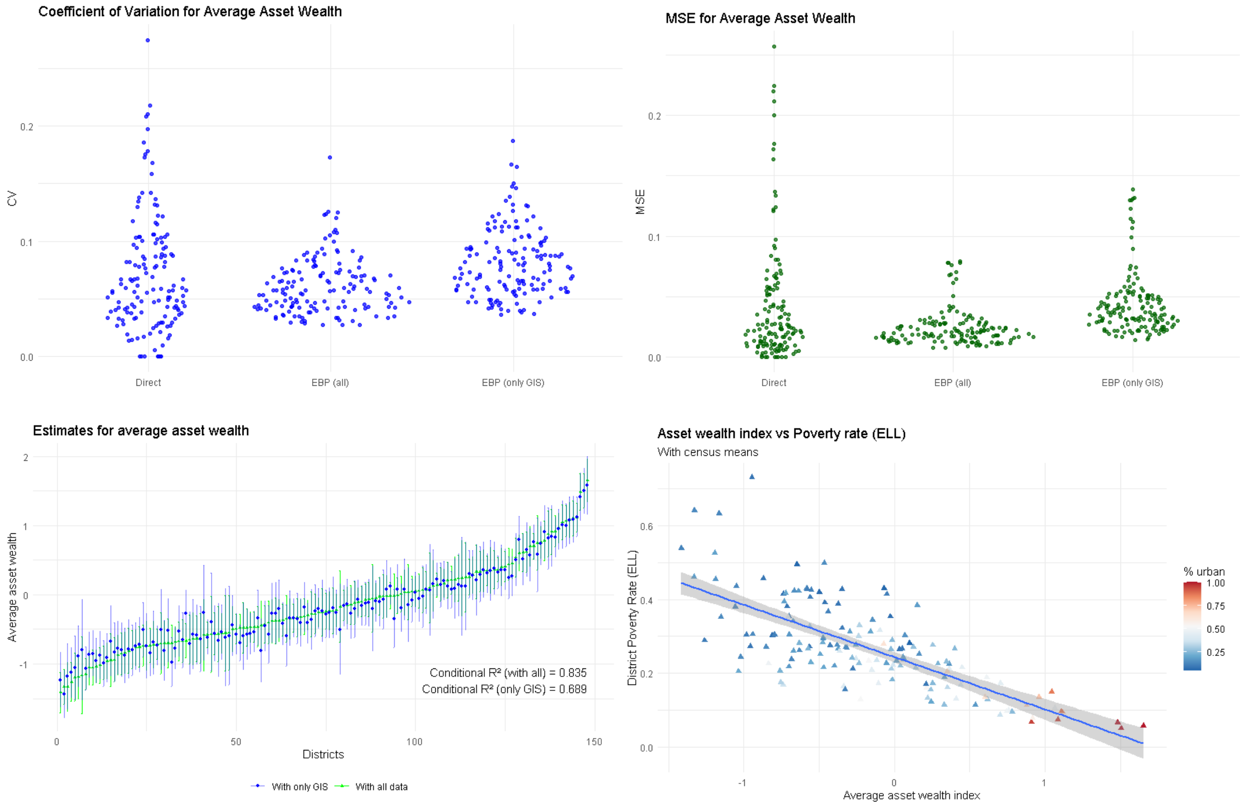

12. Application of PEBP to estimate average asset wealth index at district level. The coefficient of variation (top-left) and the MSE (top-right) for direct estimates and the PEBP estimates using all auxiliary data or only GIS data. The comparison of the estimates using different sources of auxiliary data (bottom-left) and against the monetary poverty estimates from LECS (bottom-right).

The results on Figure 12 also show high correlation between the monetary poverty rates estimated using ELL (Lao Statistics Bureau, 2023) and the average asset wealth index estimations, especially with lower variation within districts with more urban villages. The estimation of average asset wealth across years, like 2011 and 2023 when DHS data is collected in Lao PDR, could also bring insights on the trends of economic development in these areas.

Application: Estimation of stunting rates

The application of the FH and PEBP model to estimate stunting rates follows similar specifications as in the simulation exercise when estimating poverty rates[8]. However, in this case, the full population is not the total number of households but the number of children between the age of 0 and 5. The main finding of this application is that by using the proxy of average asset wealth at district level, the prediction power of the models for stunting rates increase but not enough to reduce the bias in FH or to reduce significant the variation for PEBP.

Table 2. Variation explained by the FH or PEBP model to estimate stunting rates using or not the average wealth proxy

	FH			**PEBP**	
	With proxy	**Without proxy**		**With proxy**	**Without proxy**
Adjusted R2	0.58	0.54	**Marginal**	0.34	0.29
FH R2	0.72	0.67	**Conditional**	0.39	0.37

Table 2 validates the assumption that by including the proxy of estimated average wealth at district level as auxiliary data, the explanatory power increases in both models. However, the increment is larger when considering the FH R-squared in FH or the marginal R-squared in PEBP. Contrary to the simple adjusted R-squared, the FH R-squared accounts for variation due to sampling error, which could imply that the inclusion of the proxy does increase the explanation over the variation of the outcome and not random noise (Lahiri & Suntornchost, 2015). In the case of the PEBP model, the marginal R-squared also focuses only on the variance explained by the auxiliary data, and it shows a modest increase from 0.29 to 0.34. However, the average wealth proxy is a variable measured with error by design, thus future research could test the inclusion of this known error in a FH model as suggested by Ybarra & Lohr (2008).

---

[8] Also, the auxiliary data used in models where the average wealth proxy was included, did not considered the variables used for the construction of the proxy.

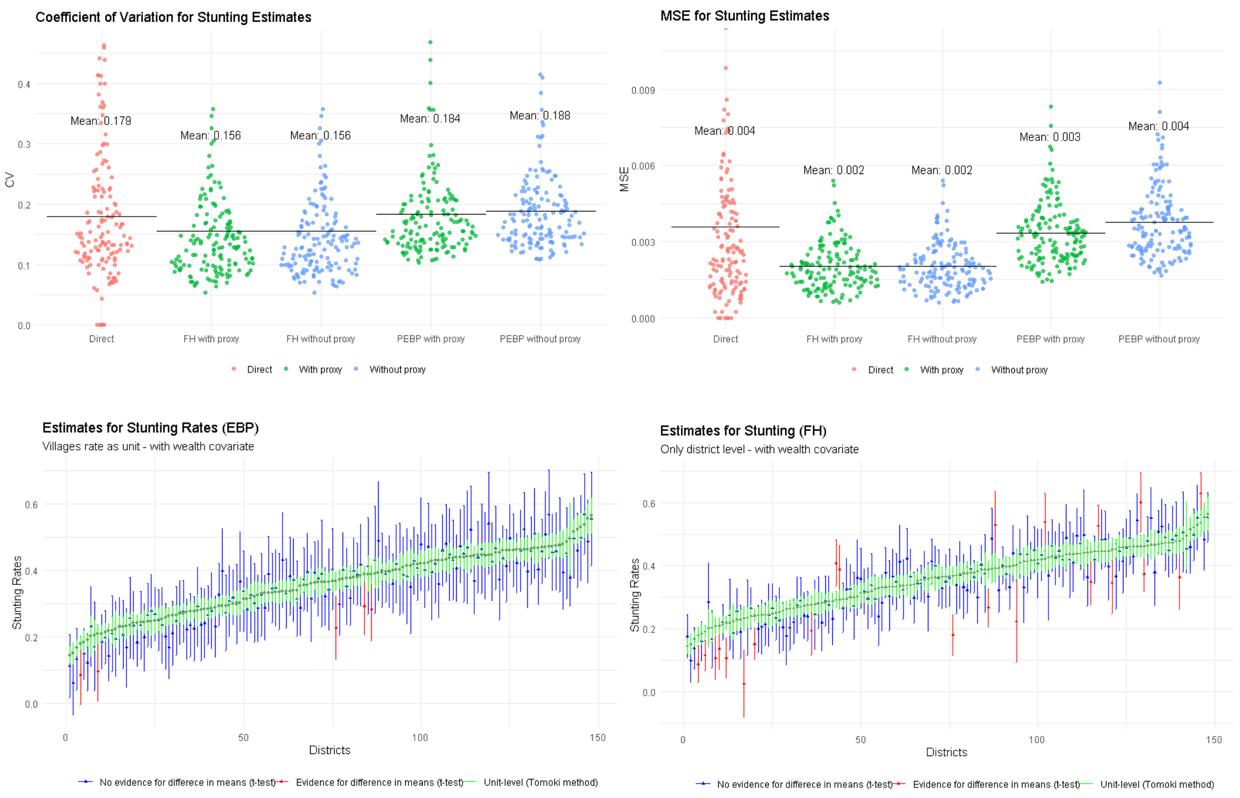

13. Application of PEBP and FH models to estimate stunting rates at district level. The coefficient of variation (top-left) and the MSE (top-right) for direct estimates and the PEBP estimates with and without the average wealth proxy as auxiliary data. The comparison using mean difference t-test of the PEBP estimates (bottom-left) and the FH (bottom-right) estimates against the estimates from unit-level using Tomoki's method (Fujii, 2005; Garcia Rojas et al., 2024).

Finally, Figure 13 illustrates the results from estimating stunting rates with FH and PEBP models. The area-level model presented lower levels regarding CV and MSE than direct and PEBP models, as expected. However, these are not affected by the inclusion of the average wealth proxy. On the other hand, when using the PEBP model, there is a slight decrease in the same parameters for validation (CV and MSE) when including the proxy. When comparing the final estimates from FH or PEBP with the estimates from a unit-level model, which uses children as the unit instead of the village (Fujii, 2005), the results appear to match in most districts. However, considering the findings from the simulation, these estimates are much noisier because of high variance and bias due to sampling errors and model specifications. A good follow-up to this exercise could be an alternative method that combines the FH and the PEBP-cluster-unit models in some way.

Conclusion

This study provides evidence that while remote sensed data can be a valuable resource for estimating welfare indicators, its effectiveness varies significantly across geographical contexts, the target estimation level and the type of indicator. In particular, the estimations based solely on GIS data tend to be more accurate in richer and urban areas. Therefore, in rural countries like Lao PDR, estimations of well-being using only remote sensed data should be used with caution. However, when the GIS data has large explanatory power and the outcomes are linear, estimation in off-census years could be achievable.

The analysis also shows that unit-level models using clusters as units, especially the PEBP model, can perform as well as an area-level model for linear indicators like average asset wealth. Nonetheless, the PEBP models offer a more balanced trade-off between bias and variance, making them a promising alternative when household-level data is unavailable. However, the precision of these estimates remains limited, particularly for nonlinear indicators. Furthermore, using average asset wealth estimates as proxies for estimation of stunting improves the explanatory power of malnutrition models.

Overall, the findings suggest that in rural countries like Lao PDR, combining remote sensed data with census aggregates and leveraging unit-level models can enhance the accuracy of welfare and malnutrition estimations. The limitations of this study are based on the simplification, and thus strong assumptions, imposed on PEBP models of null measurement error from the dependent variable, when using cluster averages, and from the independent variable, when using asset wealth estimations as a proxy for stunting rates estimation. Future research should explore hybrid modeling approaches that integrate the strengths of both area-level and unit-level methods, and account for measurement error in both outcome and auxiliary variables to improve reliability and policy relevance.

References

Asian Development Bank. (2017). *Lao PDR: Accelerating Structural Transformation for Inclusive Growth—Country Diagnostic Study* (0 ed.). Asian Development Bank. https://doi.org/10.22617/TCS179052-2

Battese, G. E., Harter, R. M., & Fuller, W. A. (1988). An Error-Components Model for Prediction of County Crop Areas Using Survey and Satellite Data. *Journal of the American Statistical Association*, *83*(401), 28–36. https://doi.org/10.1080/01621459.1988.10478561

Chen, S., Olofsson, P., Saphangthong, T., & Woodcock, C. E. (2023). Monitoring shifting cultivation in Laos with Landsat time series. *Remote Sensing of Environment*, *288*, 113507. https://doi.org/10.1016/j.rse.2023.113507

Chi, G., Fang, H., Chatterjee, S., & Blumenstock, J. E. (2022). Microestimates of wealth for all low- and middle-income countries. *Proceedings of the National Academy of Sciences*, *119*(3), e2113658119. https://doi.org/10.1073/pnas.2113658119

Corral, P., Henderson, H., & Segovia, S. (2025). Poverty mapping in the age of machine learning. *Journal of Development Economics*, *172*, 103377. https://doi.org/10.1016/j.jdeveco.2024.103377

Corral, P., Molina, I., Cojocaru, A., & Segovia, S. (2022). *Guidelines to Small Area Estimation for Poverty Mapping*. World Bank. https://doi.org/10.1596/37728

Elbers, C., Lanjouw, J. O., & Lanjouw, P. (2003). Micro-Level Estimation of Poverty and Inequality. *Econometrica*, *71*(1), 355–364. https://doi.org/10.1111/1468-0262.00399

Fay, R. E., & Herriot, R. A. (1979). Estimates of Income for Small Places: An Application of James-Stein Procedures to Census Data. *Journal of the American Statistical Association*, *74*(366a), 269–277. https://doi.org/10.1080/01621459.1979.10482505

Fujii, T. (2005). *Micro-Level Estimation of Child Malnutrition Indicators and its Application in Cambodia No. 3662, Available at SSRN:* https://ssrn.com/abstract=770946

Garcia Rojas, D. C., Appelt, J. L., Epprecht, M., Kounnavong, S., Elbers, C., Lanjouw, P. F., & Van Vliet, J. (2024). Interactions between sustainable development goals at the district level in Lao PDR. *World Development*, *178*, 106564. https://doi.org/10.1016/j.worlddev.2024.106564

Guadarrama, M., Molina, I., & Rao, J. N. K. (2018). Small area estimation of general parameters under complex sampling designs. *Computational Statistics & Data Analysis*, *121*, 20–40. https://doi.org/10.1016/j.csda.2017.11.007

Jean, N., Burke, M., Xie, M., Davis, W. M., Lobell, D. B., & Ermon, S. (2016). Combining satellite imagery and machine learning to predict poverty. *Science*, *353*(6301), 790–794. https://doi.org/10.1126/science.aaf7894

Jiang, J., & Lahiri, P. (2006). Mixed model prediction and small area estimation. *TEST*, *15*(1), 1. https://doi.org/10.1007/BF02595419

Jiang, J., Lahiri, P., & Wan, S.-M. (2002). A unified jackknife theory for empirical best prediction with M-estimation. *The Annals of Statistics*, *30*(6). https://doi.org/10.1214/aos/1043351257

Khachiyan, A., Thomas, A., Zhou, H., Hanson, G., Cloninger, A., Rosing, T., & Khandelwal, A. K. (2022). Using Neural Networks to Predict Microspatial Economic Growth. *American Economic Review: Insights*, *4*(4), 491–506. https://doi.org/10.1257/aeri.20210422

Lahiri, P., & Suntornchost, J. (2015). *Variable Selection for Linear Mixed Models with Applications in Small Area Estimation.*

Lao Statistical Bureau. (2017). *Lao PDR- Social Indicator Survey II (MICS-DHS)* [Dataset].

Lao Statistics Bureau. (2023). *Where are the poor in Lao PDR - Small Area Estimation: Province and District level Results* [Dataset].

Lee, K., & Braithwaite, J. (2022). High-resolution poverty maps in Sub-Saharan Africa. *World Development*, *159*, 106028. https://doi.org/10.1016/j.worlddev.2022.106028

Machicao, J., Ben Abbes, A., Meneguzzi, L., Corrêa, P. L. P., Specht, A., David, R., Subsol, G., Vellenich, D., Devillers, R., Stall, S., Mouquet, N., Chaumont, M., Berti-Equille, L., & Mouillot, D. (2022). Mitigation Strategies to Improve Reproducibility of Poverty Estimations From Remote Sensing Images Using Deep Learning. *Earth and Space Science*, *9*(8), e2022EA002379. https://doi.org/10.1029/2022EA002379

Molina, I., Corral, P., & Nguyen, M. (2022). Estimation of poverty and inequality in small areas: Review and discussion. *TEST*, *31*(4), 1143–1166. https://doi.org/10.1007/s11749-022-00822-1

Molina, I., Nandram, B., & Rao, J. N. K. (2014). Small area estimation of general parameters with application to poverty indicators: A hierarchical Bayes approach. *The Annals of Applied Statistics*, *8*(2). https://doi.org/10.1214/13-AOAS702

Molina, I., & Rao, J. N. K. (2010). Small area estimation of poverty indicators. *Canadian Journal of Statistics*, *38*(3), 369–385. https://doi.org/10.1002/cjs.10051

Newhouse, D. (2024). Small Area Estimation of Poverty and Wealth Using Geospatial Data: What have We Learned So Far? *Calcutta Statistical Association Bulletin*, *76*(1), 7–32. https://doi.org/10.1177/00080683231198591

Permatasari, N., & Ubaidillah, A. (2025). Small Area Estimation of poverty using remote sensing data. *Statistical Journal of the IAOS: Journal of the International Association for Official Statistics*, *41*(1), 180–190. https://doi.org/10.1177/18747655241308390

Rutstein, S., & Johnson, K. (2004). *The DHS wealth index*. http://dhsprogram.com/pubs/pdf/CR6/CR6.pdf.

Sakai, T. (2025). The probability smoothing problem: Characterizations of the Laplace method. *Mathematical Social Sciences*, *135*, 102409. https://doi.org/10.1016/j.mathsocsci.2025.102409

Sugasawa, S., & Kubokawa, T. (2017). Transforming response values in small area prediction. *Computational Statistics & Data Analysis*, *114*, 47–60. https://doi.org/10.1016/j.csda.2017.03.017

Van Der Weide, R., Blankespoor, B., Elbers, C., & Lanjouw, P. (2024). How accurate is a poverty map based on remote sensing data? An application to Malawi. *Journal of Development Economics*, *171*, 103352. https://doi.org/10.1016/j.jdeveco.2024.103352

Ybarra, L. M. R., & Lohr, S. L. (2008). *Small Area Estimation When Auxiliary Information Is Measured with Error.* http://www.jstor.org/stable/20441515

Yeh, C., Perez, A., Driscoll, A., Azzari, G., Tang, Z., Lobell, D., Ermon, S., & Burke, M. (2020). Using publicly available satellite imagery and deep learning to understand economic well-being in Africa. *Nature Communications*, *11*(1), 2583. https://doi.org/10.1038/s41467-020-16185-w