

Hybrid Fay-Herriot Random Forest for Small Area Estimation of Per Capita Expenditure in Bandung Regency

Nabillah Rahmatiah Tangke^{1*}, Asep Saefuddin¹, Anang Kurnia¹

¹Statistics and Data Science Study Program, IPB University, Indonesia

nabilatangke@apps.ipb.ac.id

ABSTRACT

Article History:

Received : 26-05-2026

Revised : 27-06-2026

Accepted : 27-06-2026

Online : 01-10-2026

Keywords:

Area Level;

EBLUP;

Fay Herriot-Random

Forest;

Lognormal

transformation;



Direct survey estimates for small areas tend to be too noisy to use on their own when only 15 to 20 households are sampled per area, which is the situation at the sub-district level in Bandung Regency. The Fay-Herriot model in Small Area Estimation (SAE) helps by borrowing strength across areas, but it leans on two assumptions that do not hold well for sub-district per capita expenditure: linearity in the auxiliary variables, and normality of the sampling errors and the area random effects. In practice, PODES indicators relate to expenditure through thresholds and interactions rather than in a straight line, and per capita expenditure at the sub-district level is right-skewed. Both are problems for the classic Fay-Herriot model, and we address them in a single model. Methodologically, it is a quantitative, model-based comparative study that aims to develop and evaluate a hybrid area-level estimator that relaxes the linearity assumption of the Fay-Herriot model and to assess its accuracy for estimating sub-district per capita expenditure. The linear fixed effect is replaced by a Random Forest, yielding a hybrid Fay-Herriot Random Forest (FH-RF), and the response is log-transformed with a Slud and Maiti style bias correction when predictions are mapped back to the original scale. The model is applied to average weekly per capita expenditure on food, beverages, and tobacco across all 31 sub-districts of Bandung Regency, with the response drawn from the 2023 Human Development Index Survey and ten auxiliary variables from PODES. Eight scenarios were compared, varying the model type, the response transformation, and the source of sampling variance. The hybrid FH-RF Log-Normal model with GVF-smoothed variance performed best, with a mean RRMSE of 7.14%, against 9.05% for the strongest classic Fay-Herriot model and 12.75% for the direct estimate, a 44.0% improvement over direct estimation and 21.1% over the best classic specification. 28 of the 31 sub-districts (90.3%) fell below 10% RRMSE, indicating that the hybrid estimator delivers consistently reliable sub-district welfare estimates for poverty targeting.



<https://doi.org/10.31764/jtam.v10i4.39931>



This is an open access article under the **CC-BY-SA** license

A. INTRODUCTION

Model-based Small Area Estimation (SAE) is the usual approach when direct survey estimates are unreliable for small areas. By borrowing strength from other areas through auxiliary variables, SAE stabilises the estimates and lowers their coefficient of variation (Rao & Molina, 2015). This idea was introduced in the seminal area-level application of Fay & Herriot (1979), who combined direct census estimates with auxiliary regression information to improve income estimates for small places. The Fay-Herriot model is the most widely used area-level model in this framework: it combines direct survey estimates with auxiliary information through a linear mixed model to produce the Empirical Best Linear Unbiased

Predictor (EBLUP) (Rao & Molina, 2016). A defining feature of SAE problems is the need to produce not only point estimates but also valid measures of uncertainty for each individual area (Pfeffermann, 2013).

The Fay-Herriot model rests on two assumptions that often fail with socioeconomic data: linearity in the auxiliary variables, and normality of both the sampling errors and the area random effects. Real relationships rarely follow a straight line. Expenditure responds to PODES indicators through thresholds and interactions that a linear regression cannot represent, and forcing a linear fit in this setting increases the risk of misspecification and unreliable uncertainty estimates. Normality is the second weak point. Expenditure data at the small-area level are typically right-skewed, which throws off the shrinkage factors and inflates the MSE. The standard remedy is to log-transform the response so it is closer to normal before fitting (Slud & Maiti, 2006; Tanur et al., 2021). Krennmair & Schmid (2022) make the broader case that machine learning offers nonlinear, nonparametric alternatives that are less vulnerable to this kind of misspecification than the usual mixed linear approach.

A number of authors have responded to these limitations with more flexible SAE methods. Parker (2024) replaces the linear component with a randomly weighted neural network inside a hierarchical Bayesian setup and shows that nonlinear models can outperform linear ones in prediction. Krennmair & Schmid (2022) developed Mixed Effects Random Forests (MERFs), where a Random Forest sits inside the mixed model, and the area random effects are preserved. Bukhari et al. (2025) pushed this further with a hybrid Fay-Herriot Random Forest (FH-RF) that swaps the linear fixed-effect component for a Random Forest, dropping the fixed distributional assumption entirely. Other machine learning methods have entered this area too, using extreme gradient boosting (Viljanen et al., 2022) and generalized random forests (Frink & Schmid, 2024). Also, related work has extended the area-level model to compositional outcomes (Esteban et al., 2020).

In the context of development data in Indonesia, several studies have applied SAE methods to estimate household consumption per capita expenditure using linear Fay-Herriot models, such as Ubaidillah et al. (2019) who developed multivariate Fay-Herriot models for household food and non-food expenditure in Central Java. Earlier, Kurnia et al. (2009) reviewed a range of SAE techniques for the Indonesian context and, in particular, extended the EBLUP to a logarithmic scale with a modified Prasad-Rao MSE to accommodate the skewed nature of expenditure data. More recently, Fauziah et al. (2025) applied a unit-level Empirical Best Predictor with a log-normal transformation to estimate per capita expenditure at the sub-district level in Garut Regency, West Java, again to address right-skewness. Closest to the present work, Bukhari et al. (2025) modified the area-level Fay-Herriot model by replacing the linear fixed effect with a Random Forest estimated from out-of-bag predictions, after initialising parameters through an EM algorithm, and applied it to per capita expenditure among agricultural and poor households. Hybrid Fay Herriot and machine learning approaches that can handle nonlinearity, however, have rarely been applied to the components of the Human Development Index (HDI).

This study develops a hybrid Fay-Herriot FH-RF model to estimate average per capita expenditure on food, beverages, and tobacco at the sub-district level in Bandung Regency. The Bandung Regency HDI Survey samples only 15 to 20 households per sub-district, too few to

yield estimates precise enough for reliable targeting, motivating a model-based SAE approach. The model relaxes the linearity assumption of the classic Fay-Herriot model: it keeps the small-area stabilization of EBLUP, but replaces the linear fixed effect with a Random Forest that can pick up nonlinear patterns in the auxiliary variables. We hope the results contribute to filling a gap in the application of hybrid SAE methods to Indonesian development data, and that the sub-district welfare map is useful to the Bandung Regency Government for targeting poverty alleviation programs in line with SDG 1.

B. METHODS

1. Data and Research Variable

This study employs two Small Area Estimation (SAE) model approaches to estimate the average per capita consumption and expenditure on food, beverages, and tobacco during the last week across 31 sub-districts in Bandung Regency in 2023. This is a quantitative study that combines a model-based estimation design with a comparative evaluation across competing estimators. It draws on secondary, area-level aggregate data, and the modeling approach is area-level Small Area Estimation: the classic Fay-Herriot EBLUP serves as the baseline, and the proposed hybrid Fay-Herriot Random Forest serves as the alternative, with both evaluated under a common cross-sectional design rather than a train-test split, since the target is simultaneous estimation for all 31 sub-districts. The response data were obtained from the Bandung Regency Human Development Index (HDI) Survey, while 10 predictor variables were sourced from Village Potential Statistics (PODES) data as area-level aggregate data, as shown in Table 1.

Table 1. Response and Auxiliary Variables Used in Per Capita Expenditure Estimation

Variable		Data Source
Average weekly consumption and expenditure on food, beverages, and tobacco	Y	IPM Survey, Bandung Regency
Proportion of urban village status	X_1	PODES 2024
Proportion of income sources from the manufacturing industry	X_2	PODES 2024
Average distance to permanent market	X_3	PODES 2024
Average distance to educational facilities	X_4	PODES 2024
Total slum households	X_5	PODES 2024
Total permanent markets	X_6	PODES 2024
Total electricity users	X_7	PODES 2024
Total poverty certificates	X_8	PODES 2024
Total landline telephones	X_9	PODES 2024
Total shops/stalls	X_{10}	PODES 2024

2. Analysis stages

Figure 1 gives the full workflow. The response is the average weekly per capita expenditure on food, beverages, and tobacco from the 2023 Bandung Regency HDI Survey, and ten auxiliary variables come from PODES 2024. We start with exploratory checks on both variable distributions, the correlation matrix, and the response turns out to be clearly right-skewed, which by itself is enough reason to carry a log-transformed copy of the response into the

modelling step. The sampling variance is handled in two parallel tracks: one uses the raw direct survey variance, the other smooths that variance with a Generalized Variance Function (GVF) before passing it to the Fay-Herriot model. Four models are then fit on each of the two datasets. Two are classic Fay-Herriot with EBLUP-REML, one on the original scale and one on the log scale; the other two are the hybrid FH-RF on the same two scales. Eight runs altogether. The log-scale predictions are back-transformed to the original expenditure scale using the Slud and Maiti (2006) bias correction.

The RF component was fit using the Random Forest package in R. Each forest was grown with $n_{tree} = 100$ regression trees, and the number of candidate variables sampled at each split was set to $m_{try} = 4$, with the default minimum node size for regression. Because the design is cross-sectional and the estimation target covers all 31 sub-districts simultaneously, the $f(X)$ was obtained from the out-of-bag (OOB) predictions rather than from an external train-test split, which would discard areas that must themselves be estimated. The mean squared error of the FH-RF estimator was obtained through a parametric bootstrap replicate. Mean RRMSE was adopted as the primary evaluation indicator because it is scale-relative and therefore comparable across sub-districts of differing expenditure levels, while mean MSE and mean RMSE are reported alongside it to show absolute error magnitude, as shown in Figure 1.




Figure 1. Research workflow for hybrid Fay-Herriot Random Forest small area estimation

3. Small Area Estimation

In surveys, sometimes we want to obtain specific information for small groups in the population. Such groups are called "small domain" because the sample size in surveys is usually very small or even zero, making the estimation unstable (Jiang & Rao, 2020; Rao & Molina, 2015). The area level is an estimation model that focuses parameter estimation on small area units, such as sub-districts or villages, based on information available at the aggregate area level. The Fay-Herriot model is one of the most popular area-level models for SAE (Fay & Herriot, 1979). This model assumes that direct estimates from surveys can be modeled as observations from area parameters that follow a normal distribution with known sampling variance. The Fay-Herriot model estimates variance parameters directly from the data (Datta & Ghosh, 2012)

$$y_i = \theta_i + e_i, i = 1, \dots, m \quad (1)$$

$$\theta_i = \mathbf{X}_i' \boldsymbol{\beta} + v_i \quad (2)$$

y_i is the direct estimator for the i -th area, θ_i is the target parameter; in addition, there is e_i which is the error/residual and $X_i = (X_1, X_2, \dots, X_{ip})$: auxiliary variables and v_i which is the random effect. Variables v_i and e_i independent with $v_i \sim N(0, \sigma_v^2)$ and $e_i \sim N(0, \sigma_{e_i}^2)$ with $i=1,2,\dots,k$. $\boldsymbol{\beta}$ and $\sigma_{e_i}^2$ are usually assumed unknown while σ_v^2 is assumed known (Kurnia et al., 2009).

4. Empirical Best Linear Unbiased Prediction (EBLUP)

EBLUP (Empirical Best Linear Unbiased Predictor) is an estimator for estimating random effects in a linear mixed model, such as the Fay-Herriot model, while the fixed effect will be estimated with weighted least squares (WLS). Diffey et al. (2017) Further developed an improved REML parameter-expanded EM algorithm for linear mixed models that ensures monotonic convergence and parameter updates remaining within the parameter space. EBLUP estimation uses the Fisher scoring algorithm based on REML (*Restricted Maximum Likelihood*) to estimate variance parameters. With this REML approach, the estimated variance value (denoted as δ). The value of δ is then used as a substitute for the value of δ in the BLUP formula, thus producing the empirical BLUP estimator, known as EBLUP, for estimating random effect parameters with the equation from Rao and Molina (2015) as follows.

$$\hat{\mu}_i^H(\sigma_v^2) = \mathbf{X}_i' \hat{\boldsymbol{\beta}} + \hat{\gamma}_i(y_i - \mathbf{X}_i' \hat{\boldsymbol{\beta}}), \quad (3)$$

$$\hat{\gamma}_i = \sigma_v^2 / (\sigma_v^2 + \sigma_{e_i}^2), = 1, \dots, m \quad (4)$$

The description of equation 2.3, where $\hat{\mu}_i^H(\sigma_v^2)$ is the EBLUP for the i -th small area, X_i is the covariate vector for the i -th area, $\hat{\boldsymbol{\beta}}$ is the regression coefficient estimated from the model, y_i is the direct estimate from area i , for example, from a survey, $\hat{\gamma}_i$ is the shrinkage factor, indicating how much weight is given to the direct estimate, σ_v^2 is the variance of the random effect between areas, and $\sigma_{e_i}^2$ for the sampling variance of area i .

Because the variance components are estimated, the MSE of the EBLUP exceeds that of the BLUP and is commonly approximated through a Taylor expansion as $g_{1i}(\hat{\sigma}_v^2) + g_{2i}(\hat{\sigma}_v^2) +$

$2g_{3i}(\hat{\sigma}_v^2)$, where g_1 is the leading prediction error, g_2 accounts for uncertainty in estimating the regression coefficients, and g_3 reflects variability from estimating the variance components; bias correction yields the standard analytic estimator under REML or method-of-moments estimation (Haslett & Welsh, 2019; Morales et al., 2021). Hirose (2019) pointed out that the choice of variance component estimator matters here: it determines whether the MSE estimator stays second-order unbiased and strictly positive. This is not a theoretical curiosity in small-area work. REML estimates of the random effect variance can collapse to zero when sample sizes are small, and when they do, the result is over-shrinkage and unreasonable MSE values.

5. EM Steps

Estimation of the variance components in the hybrid FH-RF model is performed using the Expectation-Maximization (EM) algorithm, which iteratively separates the fixed and random components until convergence. The El Leithy et al. (2016) paper focuses on EM for variance components in general mixed linear models (specifically MREML via EM), and it does confirm that EM always yields non-negative estimates as long as initial values are positive. The accuracy of mean squared error estimation through parametric bootstrap in the context of small area prediction has been further examined. The steps of the EM algorithm in the context of the Fay-Herriot model, as developed by Luis Avila-Valdez et al. (2020) are described below.

- Initialize the initial values of the regression parameter $\beta_{(0)}$, area random effect variance $\hat{\sigma}_{(0)}^2$, and set the initial iteration $r = 0$.
- E-step (expectation step), compute the estimated area of the random effect $\hat{v}_{(r)}$, using the following BLUP formula.

$$\hat{v}^{(r+1)} = \left(\frac{1}{\hat{\sigma}_{(0)}^2} \mathbf{I}_m + \mathbf{S}^{-1} \right)^{-1} \mathbf{S}^{-1} (\hat{\boldsymbol{\theta}} - \mathbf{X} \boldsymbol{\beta}_{(0)}) \quad (5)$$

- M-step (maximization step), update the estimated regression parameter and random effect variance that were previously initialized, using ML (maximum likelihood) as follows. Estimated regression parameter. $\hat{\beta}^{(r+1)} = (\mathbf{X}^T \mathbf{S}^{-1} \mathbf{X})^{-1} \mathbf{X}^T \mathbf{S}^{-1} (\hat{\boldsymbol{\theta}} - \hat{\mathbf{v}}^{(r+1)})$, and estimated random effect variance; $\hat{\sigma}^{(r+1)2} = \frac{1}{m} \left(\hat{\mathbf{v}}^{(r+1)T} \hat{\mathbf{v}}^{(r+1)} + \text{tr} \left(\frac{1}{\hat{\sigma}_{(0)}^2} \mathbf{I}_m + \mathbf{S}^{-1} \right)^{-1} \right)$.
- Repeat the E and M steps until the parameter values converge, where the change is less than 10^{-5} .

6. Fay Herriot Random Forest (FH-RF)

FH-RF is an extension of the random forest method specifically designed to handle clustered data. FH-RF is also capable of performing estimation without relying on strict distributional assumptions, which often limits other approaches in SAE (Krennmair & Schmid, 2022). The Random Forest algorithm Breiman (2001) is an ensemble tree-based method known for its robust predictive performance and ability to handle complex nonlinear interactions without explicit model assumptions (Schonlau & Zou, 2020). The FH-RF model in mathematical form is as follows:

$$y_i = f(X_i) + Z_i v_i + e_i, i = 1, \dots, m \quad (6)$$

$f(X_i)$ is the fixed component predicted using RF that models nonlinearity among variables, $Z_i v_i$ is the random effect for cluster i , where Z_i is the matrix that maps random effects v_i .

7. Lognormal Distribution for Fay Herriot

In probability theory, the lognormal distribution is a continuous probability distribution of a random variable whose logarithmic value follows a normal distribution. If the random variable X follows a lognormal distribution, then $Y = \ln(X)$ follows a normal distribution. A positive random variable X has a lognormal distribution if the logarithm of X is normally distributed ($\ln(X) \sim N(\mu, \sigma^2)$). The mean and variance can be derived as follows (Chandra et al., 2018; Ratnaparkhi, 2014).

$$E[X] = e^{\mu + \frac{1}{2}\sigma^2}, \quad (7)$$

$$E[X^2] = e^{2\mu + 2\sigma^2}, \quad (8)$$

$$Var[X] = e^{2\mu + \sigma^2}(e^{\sigma^2} - 1) \quad (9)$$

$$SD[X] = e^{\mu + \frac{1}{2}\sigma^2} \sqrt{e^{\sigma^2} - 1} \quad (10)$$

When the arithmetic mean and arithmetic variance are known, the parameters μ and σ can be obtained as follows.

$$\mu = \ln\left(\frac{E[X]^2}{\sqrt{Var[X] + E[X]^2}}\right) \quad (11)$$

$$\sigma^2 = \ln\left(1 + \frac{Var[X]}{E[X]^2}\right) \quad (12)$$

In the context of the Fay-Herriot model, the logarithmic transformation is applied to the direct estimate so that $y_i = \log(E_i)$ where E_i is the direct estimate on the original scale. The target small area parameter is then defined on the exponential scale as $\theta_i^* = \exp(\theta_i) = \exp(X_i' \beta + v_i)$. Slud & Maiti (2006) developing a bias-corrected EBLUP estimator on the original scale formulated as follows.

$$\theta_i^* = \exp\{X_i' \hat{\beta} + \hat{\gamma}_i(y_i - X_i' \hat{\beta}) + \frac{1}{2} \hat{\sigma}^2(1 - \hat{\gamma}_i)\} \quad (13)$$

The above equation has y_i which is the direct estimate on the logarithmic scale. In addition, there is $\frac{1}{2} \hat{\sigma}^2(1 - \hat{\gamma}_i)$ is the correction factor for bias that is caused by the back-transformation of the log-normal.

8. Generalized Variance Function

In complex surveys, direct variance estimates for small areas are often highly unstable due to limited sample sizes. Generalized Variance Functions (GVFs) address this problem by modeling the sampling variance as a smooth parametric function of the estimate's magnitude, thereby pooling information across domains to produce more stable variance estimates (Valliant, 1987; Wolter, 2007). The functional form of a GVF is guided by the variance-mean law $V(\mu) \propto \mu^k$, where k is the variance elasticity (Siems & Münnich, 2025). Among the canonical specifications reviewed by Wolter (2007), the log-log model is adopted in this study:

$$\log \hat{V} = \beta_0 + \beta_1 \log \hat{\theta} + \varepsilon \quad (14)$$

where n_i is the sample size of the i -th area and $\varepsilon_i \sim N(0, \tau^2)$. Following You and Hidioglou (2023), and recover the smoothed variance by exponentiating the fitted value with a bias correction: $\hat{V}_i = \exp(\hat{\beta}_0 + \hat{\beta}_1 \log n_i) \times \exp\left(\frac{\hat{\tau}^2}{2}\right)$. This specification is chosen because: (1) expenditure data follow a lognormal-like distribution for which the log-log form is theoretically appropriate (Wolter, 2007); (2) it structurally prevents negative variance estimates, which is critical at sample sizes of 15–20 households Siems & Münnich (2025) and (3) it requires only sample size as an auxiliary variable, simplifying implementation. The resulting smoothed variances are treated as known in the Fay-Herriot sampling model, following standard practice in small area estimation (Rao & Molina, 2015).

C. RESULT AND DISCUSSION

1. Data Exploration

Prior to model construction, an exploratory analysis was conducted to examine the distributional properties of both the response and predictor variables, evaluate inter-variable correlations, and assess multicollinearity. These diagnostic steps informed key modeling decisions, including the application of a logarithmic transformation to the response variable and the selection of an appropriate variance specification, as shown in Figure 2.




Figure 2. Distribution of direct estimates of per capita expenditure before (left) and after (right) lognormal transformation.

The distribution characteristics of direct estimation data for average per capita expenditure at the sub-district level are visualized in the left histogram (Figure 2). The distribution clearly shows positive skewness, where data concentration is in areas with low to moderate expenditure levels, while a few areas have extremely high expenditure values. Within the *Small Area Estimation* (SAE) framework, the use of asymmetric response variables on the original scale potentially violates basic modeling assumptions and produces inefficient estimators, especially in areas with very limited sample sizes.

As a form of statistical mitigation and model assumption fulfillment, a lognormal distribution transformation was applied to the response variable. The results of this transformation are presented in the right histogram Figure. 2. Visually, the logarithmic transformation proved effective in pulling extreme values in the right tail so that the data distribution shifts to resemble a more symmetrical normal curve. This step is a crucial prerequisite in standard Fay-Herriot modeling, which bases its weighting on the assumption that the sampling error component (*sampling error*) and area random effect (*area random effect*) are asymptotically normally distributed, as shown in Figure 3.



Figure 3. Histograms and density curves of the ten PODES auxiliary variables.

Figure 3. presents histograms along with density curves for each predictor variable used in the study. This visualization aims to identify distribution patterns, the presence of outliers, and the degree of skewness (*skewness*) in each variable. The skewed distribution patterns and the presence of outliers in several variables are important considerations in the selection of estimation methods. The classic Fay-Herriot model that assumes a linear relationship can be affected by outliers and highly skewed distributions, so hybrid approaches such as Fay-Herriot Random Forest, which are more robust to non-normal distribution patterns in predictors, have the potential to provide better estimates.




Figure 4. Pearson correlation matrix among the response and auxiliary variables.

Figure 4. presents the Pearson correlation matrix among all variables used in the study. The response variable (Y) is the average per capita expenditure on food, beverages, and tobacco during the last week, while the ten predictor variables sourced from PODES data are listed in the figure caption. Several predictor variables exhibit moderate positive correlations with the response, notably the proportion of urban village status, total electricity users, and total shops or stalls, suggesting that sub-districts with more developed economic infrastructure tend to have higher per capita expenditure. Conversely, variables such as the average distance to permanent markets show a weak negative correlation with the response, indicating that greater remoteness is associated with lower expenditure levels. Among the predictor variables themselves, some pairs display moderate intercorrelations, for instance, total electricity users and total shops or stalls, which is expected given that both reflect the degree of local economic activity. However, no pairwise correlation is high enough to indicate severe collinearity.

2. Comparison of Estimation and RRMSE

The selection of the most optimal small area estimation model is not only based on meeting data assumptions but must be proven through performance evaluation metrics. In this study, various estimation scenarios were constructed, ranging from direct estimation and the standard Fay-Herriot (FH) model to the Hybrid FH-RF model, using both direct variance and the Generalized Variance Function (GVF). A summary comparison of error values from all model scenarios is presented in Table 2.

Table 2. Comparison of Mean MSE, RMSE, and RRMSE across eight Model Scenarios and Direct Estimate.

Model Scenario	Mean MSE ^a	Mean RMSE ^b	Mean RRMSE ^c
Direct Survey	4.7691	66206	12.75
FH Standard (Direct Var)	2.5205	49306	9.77
FH Standard (GVF Var)	2.0903	45203	9.05
FH Log-Normal (Direct Var)	3.0845	54232	10.45
FH Log-Normal (GVF Var)	2.812	52367	10.14
Hybrid FH-RF Standard (Direct Var)	2.1902	43133	8.25
Hybrid FH-RF Standard (GVF Var)	2.134	42391	8.1
Hybrid FH-RF Log-Normal (Direct Var)	2.028	41789	8.07
Hybrid FH-RF Log-Normal (GVF Var)	1.472	36864	7.14

Each value represents the mean computed across all 31 sub-districts for a given model: ^a Mean Mean Squared Error (MSE) across sub-districts; ^b Mean Root Mean Squared Error (RMSE) across sub-districts; ^c Mean Relative Root Mean Squared Error (RRMSE, %) across sub-districts

Table 2 summarizes the performance of eight estimation scenarios using Mean MSE, Mean RMSE, and Mean RRMSE as evaluation metrics. Mean RRMSE is adopted as the primary criterion due to its scale-relative nature, which enables fair comparison across scenarios. Direct estimation gave the highest RRMSE at 12.75%, which is what you would expect when sub-district sample sizes are only 15 to 20 households. The classic Fay-Herriot model brought this down considerably: Standard FH reached 9.77% with the direct sampling variance and 9.05% once that variance was smoothed with the GVF. The improvement with GVF confirms that smoothed sampling variances translate into better shrinkage. The FH Log-Normal specification went the other way, slightly worse at 10.45% and 10.14%. The most likely explanation is that the bias correction needed when back-transforming from the log scale amplifies estimation error when the area random effect variance is large.

The hybrid FH-RF models outperformed every classic FH variant we tested. Their RRMSE ranged from 8.25% down to 7.14%, which is the kind of gap that only makes sense if the auxiliary expenditure relationship is genuinely nonlinear; this is the same mechanism reported by Krennmair and Schmid (2022) for mixed-effects random forests and by Parker (2024) for a nonlinear neural-network Fay-Herriot model, both of which gained over their linear counterparts when the covariate relationship departed from linearity. The best scenario, hybrid FH-RF Log-Normal with GVF variance, sat at 7.14%, that 44.0% drop from direct estimation and 21.1% from the best classic model. Three model elements contribute to this result. The log transformation stabilises the variance, the GVF smooths the unstable sampling variances, and the Random Forest absorbs the nonlinear patterns in the covariates, as shown in Figure 5.




Figure 5. Estimated per capita expenditure across 31 sub-districts under eight model scenarios

Beyond the numerical summaries in Table 2, the estimation comparison line plot across all 31 sub-districts (Figure 5) offers a more intuitive picture of how each model behaves in practice. The most striking feature of the plot is the overall agreement in direction: sub-districts that rank high under one model also rank high under the others, and the same holds for low-expenditure areas. The fact that the model lines move together is itself useful information: it means no method is producing a wildly different picture of the expenditure landscape. The north-south pattern of higher spending in the urbanised north, lower in the agricultural south, comes through under every method. Where the models part ways is in how much they swing from one sub-district to the next. The classic Fay-Herriot lines (blue) jump noticeably between sub-districts, peaking at places like Margaasih and Kutawaringin and dipping at Ibun and Arjasari. With only 15 to 20 households per sub-district, the linear mixed model has limited room to smooth this kind of noise. The hybrid FH-RF lines (green and red) are calmer. They still follow the same broad peaks and valleys, but the extremes are pulled in, which makes more sense for neighboring sub-districts that look similar on the ground. This smoothing behavior is precisely what one would expect from a model whose Random Forest component can exploit complex, nonlinear interactions among the ten PODES covariates to construct a better-fitting fixed-effect surface, leaving less residual variation to be absorbed by the area random effect. Among the hybrid variants, the FH-RF Log-Normal GVF line (dark red) traces the smoothest path of all, hugging the center of the bundle where the other model lines fluctuate around it. This matters for policy. When the Bandung Regency Government ranks sub-districts to target poverty programmes, those rankings should not flip from year to year because of sampling noise. The hybrid model produces rankings that are less affected by these artefacts, which makes them more useful when allocating limited resources.




Figure 6. RRMSE across 31 sub-districts under eight model scenarios

The RRMSE line plot (Figure 6) shifts attention from the estimates themselves to how precise those estimates are. The horizontal reference lines at 10% and 15% are practical benchmarks: anything below 10% is generally treated as acceptable for official small-area statistics, while anything above 15% is usually too imprecise to be useful. Looking first at the classic Fay-Herriot models, the blue and orange lines wander above and below the 10% threshold in an irregular, almost jagged pattern. Sub-districts like Margaasih, Igun, and Baleendah stand out as chronic trouble spots where the classic model lines spike well above 10%, sometimes approaching or exceeding the 15% line. These are precisely the sub-districts whose small samples and atypical expenditure profiles make them difficult to estimate with a linear model alone. The FH Log-Normal variants generally hug their standard-scale counterparts and do not buy any clear precision gain. In several sub-districts, the orange line actually sits a bit above the blue, which is the back-transformation penalty showing up here too.

The hybrid FH-RF Log-Normal with GVF variance (dark red) looks completely different from the rest. It hugs the bottom of the plot across nearly the entire horizontal axis, staying comfortably below 10% for 28 out of 31 sub-districts. Even in the handful of areas where it edges upward, it never approaches the 15% threshold. This uniformity is arguably as important as the average RRMSE value itself: a model that is accurate on average but wildly imprecise in a few areas may still mislead local policymakers, whereas the hybrid model delivers consistently reliable estimates everywhere. One anomaly worth discussing in detail is the behavior of the FH-RF Standard with direct variance, shown as a dashed green line. In most sub-districts, it performs well, but at Arjasari, Igun, and Nagreg, it spikes dramatically, exceeding even the classic Fay-Herriot models by a wide margin. This instability is traced to the parametric bootstrap procedure used to estimate MSE on the original, untransformed scale: when the Random Forest residuals are heavy-tailed, the bootstrap can generate extreme replicates that inflate the MSE estimate for individual areas. The log-normal transformation effectively neutralizes this problem by compressing the scale before bootstrap resampling, which explains why the FH-RF Log-Normal variants show no such spikes. This finding carries a practical lesson: while the hybrid framework is powerful, combining it with a logarithmic

transformation is essential for numerical stability, particularly in datasets where the response variable exhibits the kind of right-skewed, heavy-tailed distribution that characterizes expenditure data.

The Random Forest is not used here as a stand-alone predictor of expenditure, but as a component that replaces the linear fixed effect $X'\beta$ so that nonlinearity in the auxiliary variables can be absorbed into $f(X)$. Variable importance is therefore presented not as a ranking of welfare determinants, but as a diagnostic of the mechanism behind the accuracy gain. It was measured by the permutation-based percentage increase in mean squared error (%IncMSE) for the $f(X)$ component of the winning FH-RF Log-Normal model with GVF-smoothed variance, and is reported for the ten PODES predictors in Figure 7.




Figure 7. Permutation importance (%IncMSE) of the ten PODES predictors in the Random Forest fixed-effect component of the FH-RF Log-Normal model with GVF.

The structure is concentrated: the total number of shops or stalls (%IncMSE = 2.6) and the number of permanent markets (1.9) dominate, followed by the number of slum households (0.6) and the proportion of urban villages (0.5), while the remaining predictors are near zero or slightly negative. The two leading variables are direct proxies for local commercial activity, which is theoretically consistent with their link to per capita expenditure. The near-zero and slightly negative values are expected at a sample of only 31 areas with correlated covariates, so importance should be read as a qualitative diagnostic rather than a precise ranking. Overall, the flexibility of the FH-RF is concentrated in a small set of commercial-activity predictors, and this flexible fixed-effect surface, combined with the log-normal transformation and the GVF-smoothed sampling variance, is what lowers RRMSE relative to the linear specification. This is consistent with Bukhari et al. (2025), who likewise found that a RF improved an area-level Fay-Herriot model for per capita expenditure precisely where the covariate relationship was nonlinear, while offering little gain where it was not.

Taken together, the comparisons above place this study within the broader literature on flexible SAE. The gains point in the same direction as Krennmair and Schmid (2022), Bukhari et al. (2025), and Parker (2024), supporting rather than contradicting their conclusions, even

though differences in metric and baseline mean the magnitudes are not directly comparable. What this study adds is the demonstration that pairing the hybrid estimator with a log-normal transformation and GVF-smoothed variance is what delivers numerical stability for the small, right-skewed samples typical of Indonesian sub-district expenditure data, a combination not jointly evaluated in the earlier work.

D. CONCLUSION AND SUGGESTIONS

This study developed and tested a hybrid Fay-Herriot Random Forest (FH-RF) model for average per capita expenditure on food, beverages, and tobacco at the sub-district level in Bandung Regency. Two main conclusions come out of the analysis. First, the linearity assumption of the classic Fay-Herriot model does not hold for these data. Replacing the linear predictor $X'\beta$ with a Random Forest function $f(X)$ lowered the mean RRMSE from 9.05% (best classic model) to 7.14%, a 21.1% improvement, confirming that the relationship between the PODES auxiliary variables and per capita expenditure is nonlinear. Second, the value of the log-normal transformation depends on the model it is paired with. Within the classic linear FH framework it worsened accuracy (RRMSE rose from 9.05% to 10.14%), because the back-transformation bias correction amplifies error when the fixed-effect fit leaves substantial residual structure. Paired with the FH-RF, where the Random Forest already absorbs most systematic variation, the same transformation produced the best result across all eight scenarios (RRMSE = 7.14%). GVF-smoothed variances also outperformed direct variances in every model family, and the best model kept RRMSE below 10% in 28 of the 31 sub-districts.

The main contribution of this study is therefore both methodological and practical. Methodologically, it demonstrates that a FH-RF, a log-normal bias correction, and GVF variance smoothing can be integrated within a single area-level Fay-Herriot pipeline to handle nonlinearity, skewness, and unstable sampling variance simultaneously, a combination not previously evaluated jointly for small samples. Practically, it delivers a consistent sub-district welfare map that gives the Bandung Regency Government rankings far less sensitive to sampling noise, supporting more reliable targeting of poverty alleviation programs in line with SDG 1.

A few limitations point to directions for future work. With only 31 areas, the Random Forest contributes a flexible but modest fit, so the accuracy gain rests substantially on the log-normal transformation and the GVF-smoothed variance, which also limits how precisely individual predictor effects can be characterised. The parametric bootstrap remained unstable for a few areas on the untransformed scale, and the evidence comes from a single regency and a single welfare indicator. Future studies should therefore replicate the FH-RF on larger sets of areas and other welfare indicators, and develop a more robust MSE estimator for the untransformed scale, to confirm that substituting a Random Forest into the Fay-Herriot model is beneficial whenever the auxiliary relationship is genuinely nonlinear and paired with a variance-stabilising transformation.

REFERENCES

- Breiman, L. (2001). *Random Forests* (R. E. Schapire, Ed.; Vol. 45). Kluwer Academic Publishers.
- Bukhari, A. S., Notodiputro, K. A., Indahwati, & Fitrianto, A. (2025). A modified Fay-Herriot model with machine learning for small area estimation of per capita expenditure among the poor and agricultural households. *BIO Web of Conferences*, 186. <https://doi.org/10.1051/bioconf/202518602012>
- Chandra, H., Aditya, K., & Kumar, S. (2018). Small-area estimation under a log-transformed area-level model. *Journal of Statistical Theory and Practice*, 12(3), 497–505. <https://doi.org/10.1080/15598608.2017.1415174>
- Datta, G. S., & Ghosh, M. (2012). Small Area Shrinkage Estimation. *Statistical Science*, 27, 95–114.
- Diffey, S. M., Smith, A. B., Welsh, A. H., & Cullis, B. R. (2017). A new REML (parameter expanded) EM algorithm for linear mixed models. *Australian and New Zealand Journal of Statistics*, 59(4), 433–448. <https://doi.org/10.1111/anzs.12208>
- El Leithy, H. A., Abdel Wahed, Z. A., & Abdallah, M. S. (2016). On non-negative estimation of variance components in mixed linear models. *Journal of Advanced Research*, 7(1), 59–68. <https://doi.org/10.1016/j.jare.2015.02.001>
- Esteban, M. D., Lombardía, M. J., López-Vizcaíno, E., Morales, D., & Pérez, A. (2020). Small area estimation of proportions under area-level compositional mixed models. *Test*, 29(3), 793–818. <https://doi.org/10.1007/s11749-019-00688-w>
- Fauziah, G., Kurnia, A., & Djuraidah, A. (2025). The Empirical Best Linear Unbiased Prediction and The Emperical Best Predictor Unit-Level Approaches in Estimating Per Capita Expenditure at the Subdistrict Level. *Scientific Journal of Informatics*, 12(2), 283–294. <https://doi.org/10.15294/sji.v12i2.25037>
- 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>
- Frink, N., & Schmid, T. (2024). *Small area estimation with generalized random forests: Estimating poverty rates in Mexico*. <http://arxiv.org/abs/2406.03861>
- Haslett, S. J., & Welsh, A. H. (2019). EBLUPs : Empirical Best Linear Unbiased Predictors . In *Wiley StatsRef: Statistics Reference Online* (pp. 1–7). Wiley. <https://doi.org/10.1002/9781118445112.stat08180>
- Hirose, M. Y. (2019). A class of general adjusted maximum likelihood methods for desirable mean squared error estimation of EBLUP under the Fay–Herriot small area model. *Journal of Statistical Planning and Inference*, 199, 302–310. <https://doi.org/10.1016/j.jspi.2018.07.006>
- Jiang, J., & Rao, J. S. (2020). Annual Review of Statistics and Its Application Robust Small Area Estimation: An Overview. *Downloaded from Wwww.Annualreviews.Org. Guest*. <https://doi.org/10.1146/annurev-statistics-031219>
- Krennmair, P., & Schmid, T. (2022). Flexible domain prediction using mixed effects random forests. *Journal of the Royal Statistical Society. Series C: Applied Statistics*, 71(5), 1865–1894. <https://doi.org/10.1111/rssc.12600>
- Kurnia, A., Notodiputro, K. A., Saefuddin, A., & Mangku, I. W. (2009). Indonesia Case Small Area Estimation. *ISSTEC 2009*, 1–7.
- Luis Avila-Valdez, J., Huerta, M., Leiva, V., & Trujillo, L. (2020). The Fay-Herriot Model in Small Area Estimation: EM Algorithm And Application To Official Data. *REVSTAT - Statistical Journal*, 18(5), 613–635. <https://doi.org/10.57805/revstat.v18i5.323>
- Morales, D., Esteban, M. D., Pérez, A., & Hobza, T. (2021). Mean Squared Error of EBLUPs. In *A Course on Small Area Estimation and Mixed Models* (pp. 209–238). Springer. https://doi.org/10.1007/978-3-030-63757-6_9
- Parker, P. A. (2024). Nonlinear Fay-Herriot Models for Small Area Estimation Using Random Weight Neural Networks. *Journal of Official Statistics*, 40(2), 317–332. <https://doi.org/10.1177/0282423X241244671>
- Pfeffermann, D. (2013). New Important Developments in Small Area Estimation. *Statistical Science*, 28(1), 40 – 68. <https://doi.org/10.1214/12-STS395>

- Rao, J. N. K., & Molina, I. (2015). *Small Area Estimation* (M. P. Couper, G. Kalton, J. N. K. Rao, Schwarz Norbert, & Skinner Christopher, Eds.; Second Edition). John Wiley & Sons, Inc. <https://onlinelibrary.wiley.com/doi/>
- Rao, J. N. K., & Molina, I. (2016). Empirical bayes and hierarchical bayes estimation of poverty measures for small areas. In *Analysis of Poverty Data by Small Area Estimation* (pp. 315–324). Wiley. <https://doi.org/10.1002/9781118814963.ch17>
- Ratnaparkhi, M. (2014). Lognormal Distribution - Basic. In *Wiley StatsRef: Statistics Reference Online*. Wiley. <https://doi.org/10.1002/9781118445112.stat05889>
- Schonlau, M., & Zou, R. Y. (2020). The random forest algorithm for statistical learning. *Stata Journal*, 20(1), 3–29. <https://doi.org/10.1177/1536867X20909688>
- Siems, I., & Münnich, R. (2025). Assessing the quality of generalized variance functions. *AStA Wirtsch Sozialstat Arch*, 19(3–4), 205–221. <https://doi.org/https://doi.org/10.1007/s11943-025-00365-5>
- Slud, E. V., & Maiti, T. (2006). Mean-squared error estimation in transformed Fay-Herriot models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 68(2), 239–257. <https://doi.org/https://doi.org/10.1111/j.1467-9868.2006.00542.x>
- Tanur, E., Kurnia, A., Notodiputro, K. A., & Soleh, A. M. (2021). Log-normal small area estimation with measurement error- A n application for estimating household consumption expenditure. *ICIC Express Letters*, 15(5), 449–456. <https://doi.org/10.24507/icicel.15.05.449>
- Ubaidillah, A., Notodiputro, K. A., Kurnia, A., & Mangku, I. W. (2019). Multivariate Fay-Herriot models for small area estimation with application to household consumption per capita expenditure in Indonesia. *Journal of Applied Statistics*, 46(15), 2845–2861. <https://doi.org/10.1080/02664763.2019.1615420>
- Valliant, R. (1987). Generalized Variance Functions in Stratified Two-Stage Sampling. *Journal of the American Statistical Association*, 82(398), 499–508. <https://doi.org/10.1080/01621459.1987.10478455>
- Viljanen, M., Meijerink, L., Zwakhals, L., & van de Kastele, J. (2022). A machine learning approach to small area estimation: predicting the health, housing and well-being of the population of Netherlands. *International Journal of Health Geographics*, 21(1). <https://doi.org/10.1186/s12942-022-00304-5>
- Wolter, K. M. (2007). *Introduction to Variance Estimation* (2nd ed.). Springer. <https://doi.org/10.1007/978-0-387-35099-8>
- You, Y., & Hidirolou, M. (2023). Application of Sampling Variance Smoothing Methods for Small Area Proportion Estimation. *Journal of Official Statistics*, 39(4), 571–590. <https://doi.org/10.2478/jos-2023-0026>