

Flexible domain prediction using mixed effects random forests

Patrick Krennmair¹ | Timo Schmid² 

¹Institute of Statistics and Econometrics,
Freie Universität Berlin, Berlin, Germany

²Institute of Statistics,
Otto-Friedrich-Universität Bamberg,
Bamberg, Germany

Correspondence

Timo Schmid, Institute of Statistics,
Otto-Friedrich-Universität Bamberg,
Bamberg, Germany.
Email: Timo.Schmid@uni-bamberg.de

Abstract

This paper promotes the use of random forests as versatile tools for estimating spatially disaggregated indicators in the presence of small area-specific sample sizes. Small area estimators are predominantly conceptualised within the regression-setting and rely on linear mixed models to account for the hierarchical structure of the survey data. In contrast, machine learning methods offer non-linear and non-parametric alternatives, combining excellent predictive performance and a reduced risk of model-misspecification. Mixed effects random forests combine advantages of regression forests with the ability to model hierarchical dependencies. This paper provides a coherent framework based on mixed effects random forests for estimating small area averages and proposes a non-parametric bootstrap estimator for assessing the uncertainty of the estimates. We illustrate advantages of our proposed methodology using Mexican income-data from the state Nuevo León. Finally, the methodology is evaluated in model-based and design-based simulations comparing the proposed methodology to traditional regression-based approaches for estimating small area averages.

KEYWORDS

mean squared error, official statistics, small area estimation, tree-based methods

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2022 The Authors. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* published by John Wiley & Sons Ltd on behalf of Royal Statistical Society.

1 | INTRODUCTION

Having accurate and detailed information on social and economic conditions, summarised by appropriate indicators, is imperative for the efficient implementation of policies. The term *detailed* is used to signify information that extends beyond aggregate levels into highly disaggregated geographical and other domains (e.g. demographic groups). The term *accurate* refers to information that is estimated with appropriate level of precision and is comparable over space and time. Simply analysing data from national sample surveys is not enough for achieving the dual target of accurate and detailed information. This is mainly due to the reduction of sample sizes as the level of required detail increases. The achievement of this dual target demands appropriate model-based methodology collectively referred to as Small Area Estimation (SAE).

SAE methods can be broadly divided in two classes: first, area-level models (Fay & Herriot, 1979) assume that only aggregated data for the survey and for the auxiliary information is available. Second, unit-level models (Battese et al., 1988)—further labelled as BHF—require access to the survey and to the auxiliary information on micro-level. A versatile extension of the BHF model is the EBP-approach by Molina and Rao (2010). The EBP is capable to estimate area-level means as well as other linear and non-linear indicators. Both classes (area-level and unit-level models) are predominantly regression-based models, where the hierarchical structure of observations is modelled by random effects. These linear mixed models (LMM) assume normality of random effects and error terms. Focusing on social and economic inequality data, the required assumptions for LMMs hardly meet empirical evidence. Jiang and Rao (2020) remind that optimality results and predictive performance in model-based SAE are inevitably connected to the validity of model assumptions. Without theoretical and practical considerations regarding improperly met assumptions, estimates are potentially biased and mean squared error (MSE) estimates unreliable.

One strategy to prevent model-failure, is the assurance of normality by transforming the dependent variable (Rojas-Perilla et al., 2020; Sugasawa & Kubokawa, 2017). For instance, Rojas-Perilla et al. (2020) generalise the EBP with a data-driven transformation on the dependent variable, such that normality assumptions can be met in transformed settings. Further details on how to obtain the most-likely transformation parameter that improves the performance of unit-level models are available in Rojas-Perilla et al. (2020) and Sugasawa and Kubokawa (2019) or from a more applied perspective in Tzavidis et al. (2018). Apart from transformation strategies, another alternative is the use of models with less restrictive (parametric) assumptions to avoid model-failure. For instance, Diallo and Rao (2018) and Graf et al. (2019) formulate the EBP under more flexible distributional assumptions. Alternatively, Chambers and Tzavidis (2006) propose an approach for estimating area-level means based on M-quantile models, which are a robust method avoiding distributional assumptions of LMMs including the formal specification of area-level random effects. Tzavidis et al. (2010) and Marchetti and Tzavidis (2021) extended this approach to allow for estimating more complex statistics, like quantiles of area-specific distribution functions and non-linear indicators. Semi- or non-parametric approaches for the estimation of area-level means were investigated among others by Opsomer et al. (2008). Opsomer et al. (2008) use penalised splines regression, treating the coefficients of spline components as additional random effects within the LMM setting.

A distinct methodological option to avoid the parametric assumptions of LMMs are the class of machine learning methods. These methods are not only limited to parametric models and 'learn' predictive relations from data, including higher order interactions between covariates, without explicit model assumptions (Hastie et al., 2009; Varian, 2014). Among the broad

class of machine learning methods, we focus on tree-based models and particularly on random forests (Breiman, 2001) because they exhibit excellent predictive performance in the presence of outliers and implicitly solve problems of model-selection (Biau & Scornet, 2016). In general, the predictive perspective of (tree-based) machine learning methods conceptually transfers straight forward to the methodology of unit-level SAE-models: survey data are used to construct a model with predictive covariates. Subsequently, auxiliary information from a supplementary data source (usually census, register or administrative data) is utilised to obtain predictions over sampled and non-sampled areas. From a machine learning perspective, the survey data serve as an implicit training set to construct a proper model, while supplementary data is used to predict final indicators. Nevertheless, Jiang and Rao (2020) claim, that results from machine learning methods in SAE are harder to be interpreted and justified by SAE-practitioners, compared to LMM-alternatives.

We aim to fill this gap by providing a consistent framework enabling a coherent use of tree-based machine learning methods in SAE. In particular, we incorporate random forests within the methodological tradition of SAE by proposing a non-linear, data-driven, and semi-parametric alternative for the estimation of area-level means by using Mixed Effects Random Forests (MERF) (Hajjem et al., 2014). We focus on the construction of area-level mean-estimates using MERFs for sampled and out-of-sample domains. Our proposed model-based estimator consists of a composite model of a structural component, accounting for hierarchical dependencies of survey data with random effects and a non-parametric random forest, which models fixed effects. In contrast to existing SAE-methods, our proposed approach assists SAE-practitioners by an automated model selection. We highlight strengths and weaknesses of random forests in the context of SAE, in comparison to existing (or 'traditional') SAE-methods using design- and model-based simulations. A distinct merit of this paper is the provision of a reliable bootstrap-scheme determining the uncertainty of area-level mean-estimates. Thus, this paper aims to contribute towards the trend of diversifying the model-toolbox for SAE-practitioners and researchers, while simultaneously respecting the methodological and structural nature of SAE.

The general idea of tree-based methods in SAE is not entirely new. Anderson et al. (2014) use district-level data from Peru to juxtapose the performance of LMM-based and tree-based methods for estimating population densities. Bilton et al. (2017) use classification trees for categorical variables to incorporate auxiliary information from administrative data to survey data on household-poverty in Nepal. For a continuous variable De Moliner and Goga (2018) estimate mean electricity consumption curves for sub-populations of households in France by using methods of LMMs and regression-based trees. McConville and Toth (2019) propose a regression-tree estimator for finite-population totals, which can be viewed as an automatically selected post-stratification estimator. Dagdoug et al. (2021) analyse theoretical properties of random forest in the context of complex survey data. Mendez (2008) provides theoretical and empirical considerations for using random forests in the context of SAE and compares their performance with 'traditional' unit-level LMMs for the estimation of area-level means. Although we share the general idea of Mendez (2008), the approach of this paper differs in several ways: first of all, we leave the random forests algorithm (Breiman, 2001) unchanged and explicitly estimate random effects to account for the hierarchical structure of the data. Secondly, the proposed framework of this paper is more flexible and potentially extendable to model more complex hierarchical dependency structures as well as spatial and temporal correlations. Additionally, the extension to other methods of machine learning is possible, such as support vector machines or gradient boosted trees.

The paper is organised as follows: Section 2 provides a methodological introduction to random forests and introduces MERFs based on Hajjem et al. (2014) as a method that effectively

amalgamates random forests and the possibility to account for hierarchical dependencies of unit-level observations. Additionally, we motivate a general unit-level mixed model, treating LMMs in SAE as special cases. In Section 2.3, we discuss the construction of area-level mean-estimates. Random forests promote the flexibility of predictive models, but their lack of distributional assumptions complicates inferences. As a result, Section 3 proposes a non-parametric bootstrap-scheme for the estimation of the area-level MSE. In Section 4, we use model-based simulations under complex settings to extensively discuss and compare the performance of the proposed method for point- and MSE-estimates. We claim MERFs to be a valid alternative to existing methods for the estimation of SAE-means. In Section 5, we use household income data of the Mexican state Nuevo León to estimate area-level averages and corresponding uncertainty estimates. We highlight modelling and robustness properties of our proposed methods. Section 5.3 proceeds with a design-based simulation, which assesses the quality of results in the application of Section 5.2. Furthermore, the design-based simulation contributes to a genuine demonstration of properties and advantages of MERFs in the context of SAE. Section 6 concludes and motivates further research.

2 | THEORY AND METHOD

In this section we propose a flexible, data-driven approach using random forests for the estimation of area-level means in the presence of unit-level survey data. The method requires a joint understanding of tree-based modelling techniques and concepts of SAE. We review the basic theory of random forest and discuss modifications to ensure their applicability to hierarchical data and subsequently to applications of SAE.

2.1 | Review of random forests

Random forests combine individual decision trees (Breiman et al., 1984) to improve their joint predictive power, while simultaneously reducing their prediction variance (Biau & Scornet, 2016; Breiman, 2001). Breiman (2001) extends his idea of Bagging (Breiman, 1996)—which combines predictions from single trees through a bootstrap and aggregation procedure—to random forests that apply bootstrap aggregation on decorrelated trees. Note that the two tuning parameters of random forests are the number of trees (controlling the number of bootstrap replications) and the number of variables to be selected as candidates for each split (controlling the degree of decorrelation). Because the forest is a combination of decorrelated trees, each aiming to minimise the prediction MSE, the optimal estimator for the random forest regression function $f()$ also minimises the point-wise MSE. The minimiser under squared error loss in the regression context, is given by the conditional mean of target variable y given the data (Efron & Hastie, 2016; Wager & Athey, 2018).

Random forests captivate with a lack of assumptions such as linearity or the distributional specification of model errors, however, observations are assumed to be independent. Applications of SAE are characterised by the use of hierarchical data. Ignoring the correlation between observations, generally results in inferior point-predictions and inferences. LMMs capture the dependencies between observations by random effects, while effects between covariates are modelled by linear fixed effects, resulting in an additive model of both terms. In the context of tree-based methods, Sela and Simonoff (2012) propose a semi-parametric mixed model consisting

of a random effects part and a fixed effects non-parametric tree-model. Hajjem et al. (2011) propose a similar approach under the label of mixed effect regression trees (MERT). As the superior performance of random forests over regression trees transfers to dependent data, Hajjem et al. (2014) replace the fixed effects part in MERTs by a random forest, leading to MERF. We scrutinise this approach and propose a general semi-parametric unit-level mixed model combining the flexibility of tree-based models with the structural advantages of LMMs in the next subsection.

2.2 | Mixed effects random forests

We assume a finite population which is divided into D disjunct areas U_i , with population sizes N_i , where $i = 1, \dots, D$ specifies the areas and $N = \sum_{i=1}^D N_i$ defines the population size. We assume to have a sample from this population. The number of sampled observations in area i is given by n_i , where the sample size is denoted by $n = \sum_{i=1}^D n_i$. In each sampled area we obtain j individual observations ranging from $1, \dots, n_i$.

We define the metric target variable for area i as a $n_i \times 1$ vector of individual observations $y_i = [y_{i1}, \dots, y_{in_i}]'$. Covariates are captured in the $n_i \times p$ matrix of $X_i = [x_{i1}, \dots, x_{in_i}]'$, where p defines the number of covariates. $Z_i = [z_{i1}, \dots, z_{in_i}]'$ defines the $n_i \times q$ matrix of area-specific random effect specifiers, where q describes the dimension of random effects. $v_i = [v_{i1}, \dots, v_{iq}]'$ is the $q \times 1$ vector of random effects for area i . $\epsilon_i = [\epsilon_{i1}, \dots, \epsilon_{in_i}]'$ is the $n_i \times 1$ vector of individual error terms. Observations between areas are assumed to be independent and v_i and ϵ_i are mutually independently normally distributed with the same variance-covariance matrix H_i for random effects of each area i and R_i for individual errors. A joint notation for all D areas is as follows:

$$\begin{aligned}
 y &= \text{col}_{1 \leq i \leq D}(y_i) = (y'_1, \dots, y'_D)', & X &= \text{col}_{1 \leq i \leq D}(X_i), \\
 Z &= \text{diag}_{1 \leq i \leq D}(Z_i), & v &= \text{col}_{1 \leq i \leq D}(v_i), & \epsilon &= \text{col}_{1 \leq i \leq D}(\epsilon_i), \\
 R &= \text{diag}_{1 \leq i \leq D}(R_i), & H &= \text{diag}_{1 \leq i \leq D}(H_i).
 \end{aligned}$$

The goal is to identify a relation $f()$ between covariates X and the target variable y , in order to predict values for non-sampled observations utilising available supplementary covariates from census or register information across areas. We state a model consisting of two major parts: a fixed part $f(X)$ and a linear part Zv capturing dependencies by random effects. In the following, $f()$ can be any parametric or non-parametric function that expresses the conditional mean of target variable y given covariates X :

$$y = f(X) + Zv + \epsilon, \tag{1}$$

where

$$\epsilon \sim N(0, R) \quad \text{and} \quad v \sim N(0, H).$$

Note that for each area i the following model holds:

$$y_i = f(X_i) + Z_i v_i + \epsilon_i. \tag{2}$$

The covariance matrix of the observations y is given by the block diagonal matrix $\text{Cov}(y) = V = \text{diag}_{1 \leq i \leq D}(V_i)$, where $V_i = Z_i H_i Z_i' + R_i$. We introduce model (1) in general terms to potentially allow for modelling of complex covariance and dependency structures. However, for the rest of the paper we assume that correlations arise only due to between-area variation, that is, $R_i = \sigma_e^2 I_{n_i}$ for all areas i . Note that the already mentioned LMM proposed by Battese et al. (1988) for estimating area-level means results as a special case of (1) by setting $f(\cdot)$ to be the linear model $f(X) = X\beta$, with regression parameters $\beta = [\beta_1, \dots, \beta_p]'$. Defining $f(\cdot)$ as a random forest, results in the MERF-approach proposed by Hajjem et al. (2014), which is the preferred specification throughout the rest of the paper.

Before we continue, we want to clarify consequences of distributional assumptions in (1) that mainly address the linear part of the model. The unit-level errors are assumed to follow $\epsilon \sim N(0, R)$. However, the assumption of normality does not affect the properties of the fixed part $f(X)$ and we do not require residuals to be normally distributed for the application of our proposed method. Nevertheless, for the random components part, we require a proper likelihood function to ensure that the adapted expectation-maximisation (EM) algorithm (see below for further details) for parameter estimates converges towards a local maximum within the parameter space. A normality-based likelihood function is exploited, as it has two important properties: firstly, it facilitates the estimation of random effects due to the existence of a closed-form solution of the integral of the Gaussian likelihood function. Secondly, the maximum likelihood estimate for the variance of unit-level errors is given by the mean of the unit-level residual sum of squares. The estimation of the random effects could be also done in a non-parametric way by using discrete mixtures (Marino et al., 2018, 2019). However, the modification towards a fully non-parametric formulation of model (1) is subject to further research.

For fitting the model (1) we use an approach reminiscent of the EM-algorithm similar to Hajjem et al. (2014). In short, the MERF-algorithm subsequently estimates (a) the forest function, assuming the random effects term to be correct and (b) estimates the random effects part, assuming the Out-of-Bag-predictions (OOB-predictions) from the forest to be correct. OOB-predictions utilise the unused observations from the construction of each forest's sub-tree (Biau & Scornet, 2016; Breiman, 2001). The proposed algorithm is as follows:

1. Initialise $b = 0$ and set random components $\hat{v}_{(0)}$ to zero.
2. Set $b = b + 1$. Update $\hat{f}(X)_{(b)}$ and $\hat{v}_{(b)}$:
 - (a) $y_{(b)}^* = y - Z\hat{v}_{(b-1)}$
 - (b) Estimate $\hat{f}(\cdot)_{(b)}$ using a random forest with dependent variable $y_{(b)}^*$ and covariates X . Note that $\hat{f}(\cdot)_{(b)}$ is the same function for all areas i .
 - (c) Get the OOB-predictions $\hat{f}(X)_{(b)}^{\text{OOB}}$.
 - (d) Fit a LMM without intercept and restricted regression coefficient of 1 for $\hat{f}(X)_{(b)}^{\text{OOB}}$:

$$y = \hat{f}(X)_{(b)}^{\text{OOB}} + Z\hat{v}_{(b)} + \epsilon.$$

- (e) Extract the variance components $\hat{\sigma}_{\epsilon, (b)}^2$ and $\hat{H}_{(b)}$ and estimate the random effects by:

$$\hat{v}_{(b)} = \hat{H}_{(b)} Z' \hat{V}_{(b)}^{-1} (y - \hat{f}(X)_{(b)}^{\text{OOB}}).$$

3. Repeat Step (2) until convergence is reached.

The convergence of the algorithm is assessed by the marginal change of the modified generalised log-likelihood (GLL) criterion:

$$GLL(f, v_i|y) = \sum_{i=1}^D ([y_i - f(X_i) - Z_i v_i]' R_i^{-1} [y_i - f(X_i) - Z_i v_i] + v_i' H_i^{-1} v_i + \log |H_i| + \log |R_i|).$$

In the linear case with $f = X\beta$, and for given variance components H and R , the maximisation of the GLL-criterion is equivalent to the solution of so-called mixed model equations (Wu & Zhang, 2006)—leading to best linear unbiased predictor (BLUP): $\hat{v} = HZ'V^{-1}(y - X\hat{\beta})$. For random forests, the corresponding estimator \hat{v} for known parameters H and R is given by:

$$\hat{v} = HZ'V^{-1}(y - \hat{f}(X)^{OOB}). \tag{3}$$

Mathematical details of the derivations are provided in Appendix A. This result is in line with Capitaine et al. (2021), claiming that \hat{v} is obtained by taking the conditional expectation given the data y and subsequently \hat{v} can thus be considered as the BLUP for the linear part of model (1).

The estimation of variance components in Step 2 (d) for $\hat{\sigma}_\epsilon^2$ and \hat{H} is obtained by taking the expectation of maximum likelihood estimators given the data. Although $\hat{\sigma}_\epsilon^2$ is a naive estimator within the discussed framework, it cannot be considered as a valid estimator for the variance σ_ϵ^2 of the unit-level errors ϵ . Breiman (2001) maintains that the sum of squared residuals from OOB-predictions are a valid estimator for the squared prediction error of new individual observations. However, as an estimator of the residual variance under the model, $\hat{\sigma}_\epsilon^2$ is positively biased, as it includes uncertainty regarding the estimation of function $\hat{f}(\cdot)$. Following Mendez and Lohr (2011) we use a bias-adjusted estimator for the residual variance σ_ϵ^2 from a random forest model (1) using a bootstrap bias-correction. The essential steps to obtain the corrected residual variance are summarised as follows:

1. Use the OOB-predictions $\hat{f}(X)^{OOB}$ from the final model $\hat{f}(\cdot)$ after convergence of the algorithm.
2. Generate B bootstrap samples $y_{(b)}^* = \hat{f}(X)^{OOB} + \epsilon_{(b)}^*$, where the values $\epsilon_{(b)}^*$ are sampled with replacement from the centred marginal residuals $\hat{\epsilon} = y - \hat{f}(X)^{OOB}$.
3. Recompute $\hat{f}(X)_{(b)}^{OOB}$ using a random forest with $y_{(b)}^*$ as dependent variable.
4. Estimate the correction-term $K(\hat{f})$ by:

$$\hat{K}(\hat{f}) = B^{-1} \sum_{b=1}^B [\hat{f}(X)^{OOB} - \hat{f}(X)_{(b)}^{OOB}]^2.$$

The bias-corrected estimator for the residual variance is then given by:

$$\hat{\sigma}_{bc,\epsilon}^2 = \hat{\sigma}_\epsilon^2 - \hat{K}(\hat{f}). \tag{4}$$

2.3 | Predicting small-area averages

The MERF-model (1) predicts the conditional mean on individual level of a metric dependent variable given unit-level auxiliary information. In the context of SAE, we are not interested in predictions on individual level, but in estimating indicators such as area-level means or area-level

totals (Rao & Molina, 2015). Thus, we assume the same structural simplifications as the LMM proposed by Battese et al. (1988) for estimating area-level means throughout the paper, that is, $q = 1$, Z is a $n_i \times D$ design-matrix of area-intercept indicators, $v = [v_1, \dots, v_D]'$ is a $D \times 1$ vector of random effects, and variance-covariance matrix for random effects simplifies to $H_i = \sigma_v^2$.

Firstly, we use the fact that random forest estimates of the fixed part $\hat{f}()$ express the conditional mean on unit-level. We calculate the mean-estimator for each area i based on available supplementary data sources (usually census or administrative data) by:

$$\tilde{f}(X_i) = \frac{1}{N_i} \sum_{j=1}^{N_i} \hat{f}(X_i) = \frac{1}{N_i} \sum_{j=1}^{N_i} \hat{f}(x_{ij}).$$

Secondly, we exploit the result (3) that \hat{v}_i is the BLUP for the linear part of the model (1). Therefore, the proposed estimator for the area-level mean $\mu = [\mu_1, \dots, \mu_D]'$ is given by:

$$\hat{\mu}_i = \tilde{f}(X_i) + Z_i \hat{v}_i \quad \text{for } i = 1, \dots, D. \quad (5)$$

In cases of non-sampled areas, the proposed estimator for the area-level mean reduces to the fixed part from the random forest:

$$\hat{\mu}_i = \tilde{f}(X_i).$$

We shortly discuss properties of our proposed estimator from Equation (5). The structural component $Z_i \hat{v}_i$ captures dependency and correlation structures by random effects and the expression $\tilde{f}(X_i)$ is the fixed effects predictor of the mean. For the special case, where $\hat{f}()$ is assumed to be the linear model $\hat{f}(X) = X\hat{\beta}$, with regression parameters $\hat{\beta} = [\hat{\beta}_1, \dots, \hat{\beta}_p]'$, the estimator for area-level means resembles the result of the EBLUP (Battese et al., 1988). If $\hat{f}()$ is a random forest, we face area-specific mean-estimates for fixed-effects from a highly flexible, data-driven and non-differentiable function. Two major tuning parameters affect the predictive performance of the random forest $\hat{f}()$, that is, the number of trees and the number of split-candidates at each node controlling the degree of decorrelation. In contrast to existing parametric and non-parametric methods in SAE, our proposed estimator from Equation (5) abstains from problems due to model-selection. Random forests implicitly perform optimised model-selection including higher-order effects or non-linear interactions. Although flexible approaches such as P-Splines (Opsomer et al., 2008) potentially determine non-linear relations in covariates, users have to explicitly specify model-variables and interactions to be interpolated a priori, resulting in a comparable paradigm of model-selection compared to standard LMMs. An additional property of $\hat{f}()$ is the capability to deal with high-dimensional covariate data, that is, cases where the number of covariates p is larger than the sample size n . This property might be exploited in the context of applications to alternative Big Data sources (Marchetti et al., 2015; Schmid et al., 2017).

3 | ESTIMATION OF UNCERTAINTY

The assessment of uncertainty of area-level indicators in SAE is crucial to analyse the quality of estimates. The area-level MSE is a conventional measure fulfilling this goal, but its calculation is

a challenging task. For instance, for the unit-level LMM with block diagonal covariance matrices (Battese et al., 1988), the exact MSE cannot be analytically derived with estimated variance components (González-Manteiga et al., 2008; Rao & Molina, 2015) and only partly analytical approximations are available (Datta & Lahiri, 2000; Prasad & Rao, 1990). An alternative to estimate uncertainty of the area-level indicators are bootstrap-schemes (Chambers & Chandra, 2013; González-Manteiga et al., 2008; Hall & Maiti, 2006). In contrast, general statistical results for inference of random forests are rare, especially in comparison to the existing theory of inference using LMMs. Nevertheless, we provide a theoretical discussion on the estimation of MSEs for in-sample area-level means in the spirit of Prasad and Rao (1990) based on Mendez (2008). Derivations can be found in the online supplementary materials (S1). The resulting analytical approximation is considered to be a complement to contextualise the quality of our proposed bootstrap MSE-estimator for in-sample areas. We discuss performance details in the model-based simulation in Section 4. An exact theoretical determination and discussion of asymptotic properties will be left to further research. The theoretical background of random forests grows, but mainly aims to quantify the uncertainty of individual predictions (Athey et al., 2019; Sexton & Laake, 2009; Wager et al., 2014; Wager & Athey, 2018; Zhang et al., 2019). The extension of recent theoretical results, such as conditions for the consistency of unit-level predictions (Scornet et al., 2015) or their asymptotic normality (Wager & Athey, 2018), towards area-level indicators is a conducive strategy.

In this paper, we propose a non-parametric random effect block (REB) bootstrap for estimating the MSE of the introduced area-level estimator given by Equation (5). We aim to capture the dependence-structure of the data as well as the uncertainty introduced by the estimation of model (1). Our bootstrap-scheme builds on the non-parametric bootstrap introduced by Chambers and Chandra (2013). The proposed REB bootstrap has two major advantages: firstly, empirical residuals depend only on the correct specification of the mean behaviour function $f(\cdot)$ of the model, thus the REB setting is lenient to specification errors regarding the covariance structure of the model. Secondly, the bootstrap within blocks ensures that the variability of residuals within each area is captured. We scale and centre the empirical residuals by the bias-corrected residual variance (4) in order to eliminate the uncertainty due to the estimation of $\hat{f}(\cdot)$ from the naive residuals. The steps of the proposed bootstrap are as follows:

1. For given $\hat{f}(\cdot)$ calculate the $n_i \times 1$ vector of marginal residuals $\hat{e}_i = y_i - \hat{f}(X_i)$ and define $\hat{e} = [\hat{e}'_1, \dots, \hat{e}'_D]'$.
2. Using the marginal residuals \hat{e} , compute level-2 residuals for each area by

$$\bar{r}_i = \frac{1}{n_i} \sum_{j=1}^{n_i} \hat{e}_i \quad \text{for } i = 1, \dots, D,$$

and $\bar{r} = [\bar{r}_1, \dots, \bar{r}_D]'$ indicates the $D \times 1$ vector of level-2 residuals.

3. To replicate the hierarchical structure we use the marginal residuals and obtain the $n_i \times 1$ vector of level-1 residuals by $\hat{r}_i = \hat{e}_i - 1_{n_i} \bar{r}_i$. The residuals $\hat{r} = [\hat{r}'_1, \dots, \hat{r}'_D]'$ are scaled to the bias-corrected variance $\hat{\sigma}_{bc,e}^2$ (4) and centred, denoted by $\hat{r}^c = [\hat{r}^{c'}_1, \dots, \hat{r}^{c'}_D]'$. The level-2 residuals \bar{r}_i are also scaled to the estimated variance $\hat{H}_i = \hat{\sigma}_v^2$ and centred, denoted by $\bar{r}^c = [\bar{r}^c_1, \dots, \bar{r}^c_D]'$.
4. For $b = 1, \dots, B$:

- (a) Sample independently with replacement from the scaled and centred level-1 and level-2 residuals:

$$r_i^{(b)} = \text{srswr}(\hat{r}_i^c, n_i) \quad \text{and} \quad \bar{r}^{(b)} = \text{srswr}(\bar{r}^c, D).$$

- (b) Calculate the bootstrap population as $y^{(b)} = \hat{f}(X) + Z\bar{r}^{(b)} + r^{(b)}$ and calculate the true bootstrap population area means $\mu_i^{(b)}$ as $\frac{1}{N_i} \sum_{j=1}^{N_i} y_{ij}^{(b)}$ for all $i = 1, \dots, D$.
- (c) For each bootstrap population draw a bootstrap sample with the same n_i as the original sample. Use the bootstrap sample to obtain estimates $\hat{f}^{(b)}()$ and $\hat{v}^{(b)}$ as discussed in Section 2.2.
- (d) Calculate area-level means following Section 2.3 by

$$\hat{\mu}_i^{(b)} = \hat{f}^{(b)}(X_i) + Z_i \hat{v}_i^{(b)}.$$

5. Using the B bootstrap samples, the MSE-estimator is obtained as follows:

$$\widehat{MSE}_i = B^{-1} \sum_{b=1}^B \left(\mu_i^{(b)} - \hat{\mu}_i^{(b)} \right)^2.$$

4 | MODEL-BASED SIMULATION

This section marks the first step in the empirical assessment of the proposed method. The model-based simulation juxtaposes point estimates for the area-level mean from the mixed effects random forest model (1) with several competitors. In particular, we study the performance of MERFs compared to the BHF-model (Battese et al., 1988), the EBP (Molina & Rao, 2010), the EBP under data-driven Box–Cox transformation (EBP-BC) by Rojas-Perilla et al. (2020) as well as the non-parametric EBLUP with P-Splines (P-SPLINES) by Opsomer et al. (2008). The BHF-model serves as an established baseline for the estimation of area-level means and the EBP and the EBP-BC conceptually build on the BHF-model. The non-parametric EBLUP by Opsomer et al. (2008) incorporate advantages of flexible, non-linear smoothing methods into existing theory for SAE based on LMMs. Differences in the performance of the EBP and the EBP-BC highlight advantages of data-driven transformations, while differences in the performance of the linear competitors and more flexible alternatives (MERF, P-SPLINES) indicate advantages of semi-parametric and non-linear modelling. Overall, we aim to show, that our proposed methodology for point and uncertainty estimates performs comparably well to ‘traditional’ SAE-methods and has comparative advantages in terms of robustness against model-failure.

The simulation-setting is characterised by a finite population U of size $N = 50,000$ with $D = 50$ disjoint areas U_1, \dots, U_D of equal size $N_i = 1000$. We generate samples under stratified random sampling, utilising the 50 small areas as stratas, resulting in a sample size of $n = \sum_{i=1}^D n_i = 1229$. The area-specific sample sizes range from 6 to 49 sampled units with a median of 21 and a mean of 25. The sample sizes are comparable to area-level sample sizes in the application in Section 5 and can thus be considered to be realistic.

We consider four scenarios denoted as *Normal*, *Interaction*, *Normal-Par*, *Interaction-Par* and repeat each scenario independently $M = 500$ times. The comparison of competing model-estimates under these four scenarios allows us to examine the performance under two

TABLE 1 Model-based simulation scenarios

Scenario	Model	x1	x2	μ	v	ε
Normal	$y = 5000 - 500x_1 - 500x_2 + v + \epsilon$	$N(\mu, 3^2)$	$N(\mu, 3^2)$	$\text{unif}(-1, 1)$	$N(0, 500^2)$	$N(0, 1000^2)$
Interaction	$y = 15,000 - 500x_1x_2 - 250x_2^2 + v + \epsilon$	$N(\mu, 4^2)$	$N(\mu, 2^2)$	$\text{unif}(-1, 1)$	$N(0, 500^2)$	$N(0, 1000^2)$
Normal-Par	$y = 5000 - 500x_1 - 500x_2 + v + \epsilon$	$N(\mu, 3^2)$	$N(\mu, 3^2)$	$\text{unif}(-1, 1)$	$N(0, 500^2)$	$\text{Par}(3, 800)$
Interaction-Par	$y = 20,000 - 500x_1x_2 - 250x_2^2 + v + \epsilon$	$N(\mu, 2^2)$	$N(\mu, 2^2)$	$\text{unif}(-1, 1)$	$N(0, 1000^2)$	$\text{Par}(3, 800)$

major dimensions of model-misspecification: Firstly, the presence of skewed data delineated by non-normal error-terms and secondly, the presence of unknown non-linear interactions between covariates. Scenario *Normal* provides a baseline under LMMs with normally distributed random effects and unit-level errors. As model-assumptions for LMMs are fully met, we aim to show, that MERFs perform comparably well to linear competitors in the reference scenario. Scenario *Interaction* shares its error-structure with *Normal*, but involves a complex model including quadratic terms and interactions. This scenario portrays advantages of semi-parametric and non-linear modelling methods protecting against model-failure. Working with inequality or income data, we often deal with skewed target variables. Thus, we use the Pareto distribution to mimic realistic income scenarios. Scenario *Normal-Par* uses the linear additive structure of LMMs and adds Pareto distributed unit-level errors. The resulting scenario, including a skewed target variable, is a classical example promoting the use of transformations assuring that assumptions of LMMs to be met. Finally, scenario *Interaction-Par* combines the two discussed dimensions of model misspecification, that is, a non-Gaussian error-structure with complex interactions between covariates. We chose this scenario to emphasise the ability of MERFs to handle both complications simultaneously. Further details on the data-generating process for each scenario are provided in Table 1.

We evaluate point estimates for the area-level mean by the relative bias (RB) and the relative root mean squared error (RRMSE). As quality-criteria for the evaluation of the MSE-estimates, we choose the relative bias of RMSE (RB-RMSE) and the relative root MSE of the RMSE:

$$\begin{aligned}
 \text{RB}_i &= \frac{1}{M} \sum_{m=1}^M \left(\frac{\hat{\mu}_i^{(m)} - \mu_i^{(m)}}{\mu_i^{(m)}} \right) \\
 \text{RRMSE}_i &= \frac{\sqrt{\frac{1}{M} \sum_{m=1}^M (\hat{\mu}_i^{(m)} - \mu_i^{(m)})^2}}{\frac{1}{M} \sum_{m=1}^M \mu_i^{(m)}} \\
 \text{RB-RMSE}_i &= \frac{\sqrt{\frac{1}{M} \sum_{m=1}^M \text{MSE}_{\text{est}_i}^{(m)} - \text{RMSE}_{\text{emp}_i}}}{\text{RMSE}_{\text{emp}_i}} \\
 \text{RRMSE-RMSE}_i &= \frac{\sqrt{\frac{1}{M} \sum_{m=1}^M \left(\sqrt{\text{MSE}_{\text{est}_i}^{(m)}} - \text{RMSE}_{\text{emp}_i} \right)^2}}{\text{RMSE}_{\text{emp}_i}},
 \end{aligned}$$

where $\hat{\mu}_i^{(m)}$ is the estimated mean in area i based on any of the methods mentioned above and $\mu_i^{(m)}$ defines the true mean for area i in simulation round m . $\text{MSE}_{\text{est}_i}^{(m)}$ is estimated by the proposed




FIGURE 1 Empirical root mean squared error comparison of point estimates for area-level averages under four scenarios [Colour figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com/doi/10.1111/rssc.12400)]

bootstrap in Section 3 and $\text{RMSE}_{\text{emp}_i} = \sqrt{\frac{1}{M} \sum_{m=1}^M (\hat{\mu}_i^{(m)} - \mu_i^{(m)})^2}$ is the empirical root MSE over M replications.

For the computational realisation of the model-based simulation, we use R (R Core Team, 2021). The BHF estimates are realised from the *sae*-package (Molina & Marhuenda, 2015) and the *emdi*-package (Kreutzmann et al., 2019) is used for the EBP as well as the EBP under the data-driven Box–Cox transformation. We implement the P-SPLINE method with the package *mgcv* (Wood, 2017). For estimating the proposed MERF-approach, we use the packages *randomForest* (Liaw & Wiener, 2002) and *lme4* (Bates et al., 2015). We monitor the convergence of algorithm introduced in Section 2.2 with a precision of $1e^{-5}$ in relative difference of the GLL-criterion and set the number of split-candidates to 1, keeping the default of 500 trees for each forest.

We start with a focus on the performance of point estimates. Figure 1 reports the empirical RMSE of each method under the four scenarios. As expected, in the *Normal* scenario, the BHF and the EBP perform on the same level and outperform the MERF estimator. The EBP with a data-driven transformation (EBP-BC) and the non-parametric EBLUP (P-SPLINES) lead to similar results compared to the BHF and EBP. This shows that the data-driven transformation and the penalised smoothing approach work as expected. A similar pattern appears in the results from the *Normal-Par* scenario, except that the EBP-BC reaches a lower overall RMSE due to its property of data-driven transformation and subsequently improved estimation under skewed data. As anticipated, a comparison of the performance of the MERF between the *Normal* and the

TABLE 2 Mean and Median of relative bias (RB) and relative root mean squared error (RRMSE) over areas for point estimates in four scenarios

	<i>Normal</i>		<i>Interaction</i>		<i>Normal-Par</i>		<i>Interaction-Par</i>	
	Median	Mean	Median	Mean	Median	Mean	Median	Mean
RB (%)								
BHF	0.087	0.131	-0.202	0.106	0.193	0.220	0.043	0.142
EBP	0.069	0.128	-0.060	0.108	0.216	0.217	0.105	0.142
EBP-BC	0.152	0.184	0.156	0.381	0.174	0.129	0.139	0.262
P-SPLINES	0.096	0.137	-0.064	0.123	0.199	0.227	0.051	0.090
MERF	0.137	0.191	0.279	0.312	0.409	0.460	0.151	0.188
RRMSE (%)								
BHF	3.830	4.090	3.770	3.870	3.600	4.100	2.800	2.950
EBP	3.850	4.100	3.750	3.870	3.600	4.120	2.830	2.950
EBP-BC	3.850	4.100	3.680	3.800	3.430	3.710	2.650	2.770
P-SPLINES	3.840	4.090	3.580	3.620	3.590	4.100	2.380	2.490
MERF	4.070	4.380	2.270	2.330	3.890	4.380	1.420	1.530

Abbreviations: BHF, Battese et al.; EBP-BC, EBP under data-driven Box-Cox transformation; MERF, mixed effects random forests.

Normal-Par scenario indicates, that the MERF exhibits robust behaviour under skewed data and subsequently regarding violations of the normal distribution of errors. LMM-based competitors match the data-generating process of fixed effects and perform accordingly, as already observed under the *Normal* scenario. For complex scenarios, that is, *Interaction* and *Interaction-Par*, point estimates of the proposed MERF outperform the SAE-methods based on LMMs. The EBP-BC performs better in terms of lower RMSE values compared to the BHF and the EBP in both interaction scenarios. The flexible approach of P-SPLINES outperforms the BHF, the EBP and the data-driven EBP-BC. However, MERFs automatically identify interactions and non-linear relations, such as the quadratic term in scenario *Interaction-Par*, which leads to a clear comparative advantage in terms of RMSE. Overall, the results from Figure 1 indicate that the MERF performs comparably well to LMMs in simple scenarios, and outperforms ‘traditional’ SAE-models in the presence of unknown non-linear relations between covariates. Additionally, the robustness against model-misspecification of MERFs holds if distributional assumptions for LMMs are not met, that is, in the presence of non-normally distributed errors and skewed data. Table 2 reports the corresponding values of RB and RRMSE for our discussed point estimates. The RB and the RRMSE from the MERF-method attest a competitively low level for all scenarios. Most interestingly, in complex scenarios (*Interaction* and *Interaction-Par*), a familiar result regarding the statistical properties of random forests appears: the RB is higher compared to the LMM-based models, but the enlarged RB is rewarded by a lower RRMSE for point estimates.

We scrutinise the performance of our proposed MSE-estimator on the four scenarios, examining whether the observed robustness against model-misspecification due to unknown complex interactions between covariates or skewed data for point estimates, also holds for our non-parametric bootstrap-scheme. For each scenario and each simulation round, we choose the parameter of bootstrap replications $B = 200$. From the comparison of RB-RMSE among the four

TABLE 3 Performance of bootstrap and analytical MSE estimators in model-based simulation: mean and median of RB-RMSE and RRMSE-RMSE over areas

	<i>Normal</i>		<i>Interaction</i>		<i>Normal-Par</i>		<i>Interaction-Par</i>	
	Median	Mean	Median	Mean	Median	Mean	Median	Mean
RB-RMSE (%)								
Bootstrap	0.319	-0.084	0.127	0.548	0.340	0.724	-0.802	0.123
Analytic	-5.700	-5.010	0.707	0.261	-4.020	-4.480	-7.500	-7.000
RRMSE-RMSE (%)								
Bootstrap	12.500	12.500	22.200	22.800	43.100	48.200	41.000	44.700
Analytic	6.130	5.930	10.400	12.200	21.300	21.400	33.600	33.500

Abbreviations: MSE, mean squared error; RB-RMSE, relative bias of RMSE; RMSE, root mean squared error; RRMSE-RMSE, relative root mean squared error-RMSE.

scenarios provided in Table 3, we infer, that the proposed non-parametric bootstrap procedure effectively handles scenarios that lead to model-misspecification in cases of (untransformed) LMMs. This is demonstrated by essentially unbiasedness in terms of mean and median values of RB-RMSE over areas of the MSE-estimator under all four scenarios: independently, whether the data generating process is characterised by complex interactions (*Interaction*), non-normal error terms (*Normal-Par*) or a combination of both problems (*Interaction-Par*). We compare the performance of our bootstrap estimator to an estimator resulting from an analytical discussion of uncertainty in the spirit of Prasad and Rao (1990), which can be found in online supplementary material (S1). The analytical approximation generally underestimates the MSE, except for the *Interaction* scenario, which substantiates the quality of the proposed bootstrap estimator. A detailed graphical comparison of the RB-RMSE between the non-parametric bootstrap and the analytical MSE estimator is provided by Figure B1 in the Appendix B.

From the results of Table 3 and the subsequent discussion, we cannot directly infer the area-wise tracking properties of the estimated RMSE against the empirical RMSE over our 500 simulation rounds. Thus, Figure 2 provides additional intuition on the quality of our proposed non-parametric MSE-bootstrap estimator. Given the tracking properties in all four scenarios, we conclude that our bootstrap MSE-estimates strongly correspond to the empirical RMSE and appear to track the domain-specific empirical RMSE more precisely than the estimates of our analytical MSE estimator from the theoretical discussion in online supplementary material (S1). Furthermore, we do not observe systematic differences between the bootstrapped and empirical MSE-estimates regarding different survey-sample sizes.

5 | APPLICATION: ESTIMATING HOUSEHOLD INCOME FOR MEXICAN MUNICIPALITIES

In this section, we discuss the performance of our proposed method in the context of a genuine SAE example. Concretely, we apply the MERF-method proposed in Section 2.2 to estimate domain-level average household income for the Mexican state Nuevo León. Section 5.1 describes the data and Section 5.2 reports results. We end our empirical analysis by an additional design-based simulation enabling a profound discussion on the quality of point and MSE-estimates in Section 5.3.

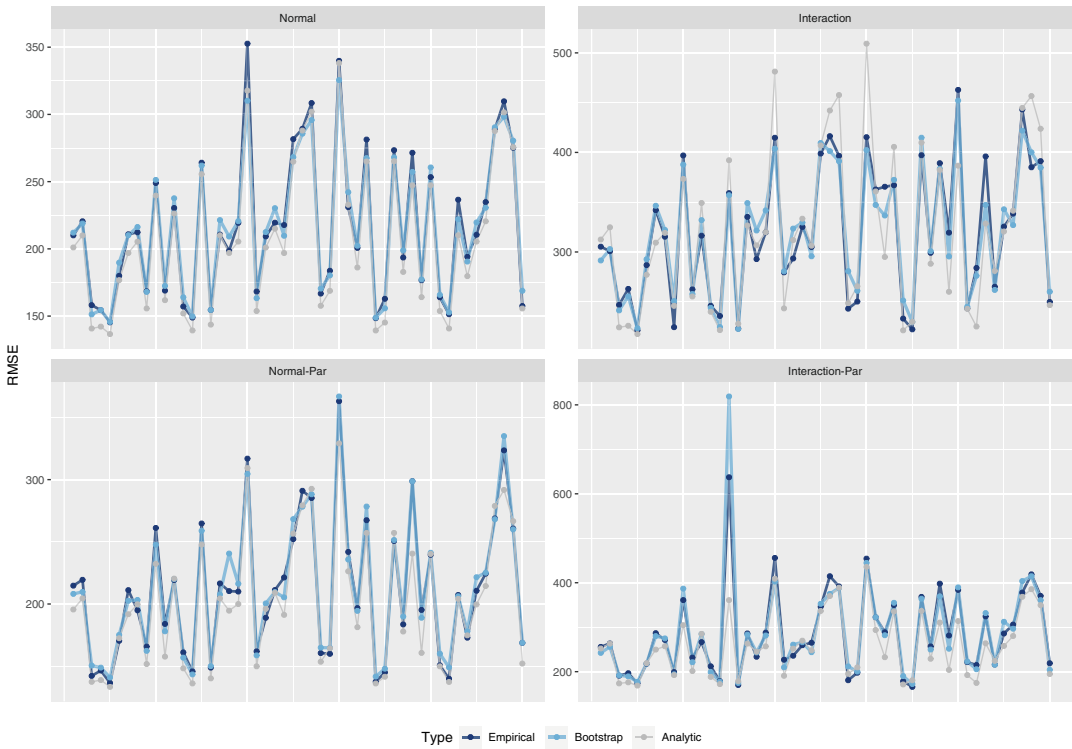


FIGURE 2 Empirical, bootstrapped and analytical area-level root mean squared errors for four scenarios [Colour figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com)]

5.1 | Data description

Income inequality in Mexico is a research topic of timeless importance, particularly regarding the effectiveness of developmental and social policies (Lambert & Park, 2019). Although the level of income inequality in Mexico is comparable to other Latin American countries, it is one of the highest among other OECD countries (OECD, 2021). Analysing average national income values is a common practice, but constitutes an inappropriate measure to monitor the efficacy of regional policy measures. Besides detailed disaggregated information, also suitable statistical methods are needed to quantify local developments. For the following application, we break down regional differences in average household per capita income of one of 32 Mexican states. Nuevo León is located in the North-East of Mexico and according to the (sub-national) Human Development Index, it is one of the most developed states of Mexico (Smits & Permanyer, 2019). Nevertheless, the distribution of individual household income in Nuevo León is unequal and thus highly skewed. For instance, the Gini-coefficient of household income is comparable to the total Gini of Mexican household disposable income from 2012 which was 0.46 (OECD, 2021).

We use data from 2010 provided by CONEVAL (Consejo Nacional de Evaluación de la Política de Desarrollo Social), combining the Mexican household income and expenditure survey (Encuesta Nacional de Ingreso y Gastos de los Hogares, ENIGH) with a sample of census micro-data by the National Institute of Statistics and Geography (Instituto Nacional de Estadística y Geografía). The data set comprises income and socio-demographic data, equally measured by variables that are part of the survey as well as the census data. The target variable for the

TABLE 4 Summary statistics on in- and out-of-sample areas: area-specific sample size of census and survey data

	Total 51		In-sample 21		Out-of-sample 30	
	Min.	First Qu.	Median	Mean	Third Qu.	Max.
Survey area sizes	5.00	14.00	27.00	68.33	79.00	342.00
Census area sizes	76.00	454.50	642.00	1075.45	872.50	5904.00

estimation of domain-level average household income in Section 5.2, is the total household per capita income (*ictpc*, measured in pesos), which is available in the survey but not in the census.

Nuevo León is divided into 51 municipalities. While the census data set in our example comprises information on 54,848 households from all 51 municipalities, the survey data includes information on 1435 households from 21 municipalities, ranging from a minimum of 5 to a maximum of 342 households with a median of 27 households. This leaves 30 administrative divisions to be out-of-sample. Table 4 provides details on sample and census data properties.

With respect to the design-based simulation in Section 5.3, we emphasise that we are in the fortunate position of having a variable that is highly correlated with the target variable *ictpc* in the application and that is available in the survey and in the census data set: the variable *inglabpc* measures earned per capita income from work. Although *inglabpc* deviates from the desired income definition for our approach—as it covers only one aspect of total household income—it is effective to evaluate our method under a design-based simulation in Section 5.3. Furthermore, the design-based simulation implicitly assesses the quality of our empirical results from Section 5.2.

Using data from Nuevo León for the estimation of domain-level income averages, is an illustrative and realistic example and imposes several challenges on the proposed method of MERFs: first of all, about 24% of households in the survey-sample are located in the capital Monterrey. Secondly, there exist more out-of-sample domains than in-sample domains. Moreover, we are confronted with households reporting zero-incomes. Our intention in choosing this challenging example for the application part in Section 5.2 and the following design-based simulation in Section 5.3 are simple: we aim to show, that our proposed approaches for point and uncertainty estimates demonstrate a valid alternative to existing SAE-methods and are predominantly applicable for cases where ‘traditional’ methods perform poorly or even fail. Additionally, we aim to provide a clear-cut presentation and empirical assessment of MERFs for SAE, which requires a transparent discussion of advantages and potential weaknesses in demanding real-world examples.

5.2 | Results and discussion

Direct estimates for average total household per capita income for Nuevo León are possible for 21 out of 51 domains. The use of model-based SAE-methods incorporating covariate census data will not only lead to estimates for the remaining out-of-sample areas, but correspondingly improve the overall quality of estimates (Tzavidis et al., 2018). As variable *ictpc* is highly skewed, we deduce potential issues of model-misspecification and suggest the use of the EBP-BC and the MERF. Given the theoretical discussion and the results in the model-based simulation in Section 4, we infer that the EBP-BC and the proposed method of MERFs for SAE effectively handle

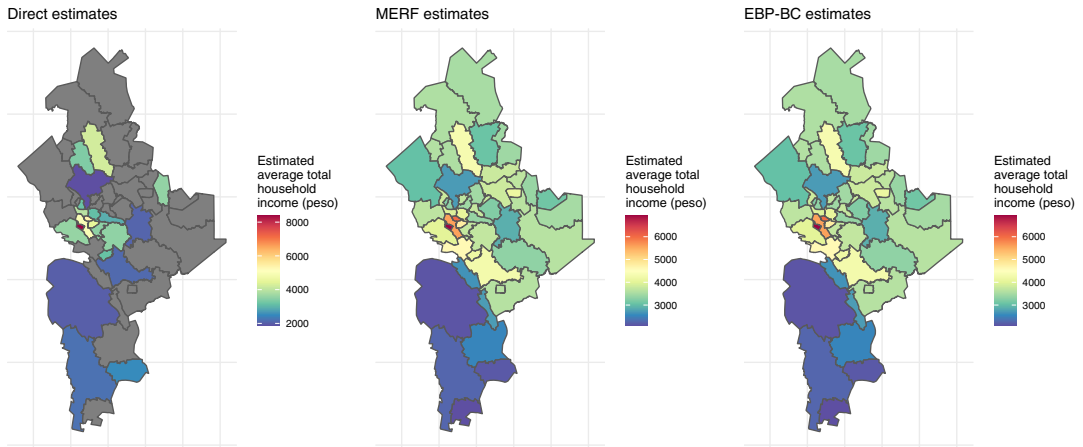


FIGURE 3 Estimated average total household per capita income *itpc* for the state Nuevo León based on three different estimation methods [Colour figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com/doi/10.1111/rssc.12400)]

non-normally distributed data. Moreover, we are particularly interested in differences between these two diverse SAE-models in the context of real-world applications. We use the design-based simulation in Section 5.3 to extend our methodological discussion towards all methods discussed in the model-based simulation in Section 4.

Figure 3 maps results from direct estimates, the MERF and the EBP-BC. Obviously, the model-based estimates from the MERF and the EBP-BC expand the perspective of regional disparities in average total household income per capita towards non-sampled regions. Furthermore, we identify three distinct clusters of income levels from our results in Figure 3 that had not been observable from the mapping of direct estimates: a low income cluster in the South of Nuevo León, a very high income cluster in the metropolitan area of the capital Monterrey and a group of middle-income areas between the North and the South of the state. This finding illustrates, the potential of model-based techniques to highlight patterns of regional income disparities and enable the mapping of reliable empirical evidence. Given the information provided by the three maps, we do not report major differences between the point estimates of MERFs and the EBP-BC.

Apart from mapping empirical results of domain averages, we are mainly interested in quality criteria, such as the coefficients of variation (CV) and the details of model-specification for the EBP-BC and the MERF. To obtain estimates of variances for the direct estimates, we use the calibrated bootstrap method (Alfons & Templ, 2013) as provided in the R package *emdi* (Kreutzmann et al., 2019). For the MSE-estimates of the MERF, we rely on the non-parametric bootstrap from Section 3. For the model of the data-driven EBP-BC, we follow the approach of Rojas-Perilla et al. (2020) and use the Bayesian Information Criterion to identify valid predictors for the target variable of *itpc*. The resulting working-model includes variables determining occupation, sources of income, the socio-economic level and educational aspects of individual households. The identification of predictive covariates for MERFs highlights a conceptual difference to LMM-based methods. Due to the properties of the random forest algorithm (Breiman, 2001), random forests perform an implicit variable selection. The selected model for fixed effects in our case is characterised by an R-squared of about 0.47%.

The dilemma between predictive precision and the interpretability of random forest models can be mitigated by concepts such as variable importance plots (Greenwell et al., 2020) (Figure 4) or an analysis of partial dependence for influential covariates (Greenwell, 2017) (Figure 5).




FIGURE 4 Variable importance in mean decrease accuracy (%IncMSE) for the ten most influential variables [Colour figure can be viewed at wileyonlinelibrary.com]

Variable importance is reported by the mean increase in individual mean squared prediction error (%IncMSE), resulting from the exclusion of the corresponding variable. Partial dependence plots depict the estimated marginal effect of a particular variable on the predicted individual target variables. From the inspection of Figure 5, we can infer whether relationships between *ictpc* and predictive variables are monotonic or more complex. Figure 4 reveals that the most important variable for the random forest model is the average relative amount of schooling (*escol_rel_hog*), followed by the availability of goods in the household (*bienes*) as well as the average years of schooling of persons (*jaesc*). Table B1 in the Appendix B provides explanations on further variables. The most influential variables are related to education, work experience and employment and household assets. Figure 5 indicates rather complex and non-linear relationships between *ictpc* and its predictive covariates except for two variables related to the number of income earners in the household (*pcpering*, *pcpering_2*).

We monitor the convergence of the proposed MERF algorithm under a precision of $1e^{-5}$ in relative difference of the GLL criterion and keep the default of 500 trees. A parameter optimisation based on fivefold cross-validation on the original survey-sample advises the use of three variables at each split for the forest. For the MSE-bootstrap procedure, we use $B = 200$.

Figure 6 reports corresponding CVs for in- and out-of-sample domains. We observe a significant improvement for in-sample CVs of the EBP-BC and the MERF compared to the CVs for direct estimates. CVs of MERFs are slightly lower in median terms than the results for the EBP-BC. However, there exists one outlying area for MERFs. Going into details, the corresponding area of




FIGURE 5 Partial dependence plots for variables ranked by %IncMSE [Colour figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com)]

General Zaragoza features no obvious data-specific irregularities, such as extremely low sample size. Nevertheless, General Zaragoza is one of the poorest regions according to our analysis. In the design-based simulation in Section 5.3, we will pay special attention to differences between MERFs and the EBP-BC regarding their ability to handle comparably extreme estimates given a broad range of relatively high and relative low income areas.

Regarding the CVs for out-of-sample areas, we discover an evident advantage of CVs from our proposed MERF approach. From the scrutiny of individual CV values it remains unclear, whether the asset of improved results from MERFs roots in superior point estimates for domain-level averages or its relatively lower MSE-estimates. Figure 7 compares direct estimates to the model-based estimates for in- and out-of-sample domains. Apparently, there exist no systematic differences between the estimates from the EBP-BC and the MERF. Thus, it appears as if the variance of MERF predictions is generally lower. This conjecture is in line with the theoretical properties of random forests (Biau & Scornet, 2016; Breiman, 2001). In-sample areas in Figure 7 are sorted by survey-sample sizes. In comparison to the direct estimates, predicted averages of the EBP-BC as well as of the MERF seem less extreme. The obvious irregularity in terms of high-income, is a distinct part of the Monterrey metropolitan area: San Pedro Garza García registers several headquarters of national and international corporations. This economic peculiarity apparently transfers to the income of its citizenship. Figure 3 underlines the existence of an apparent high-income cluster in this region. Overall, it is interesting to observe how reliable estimates on fine spatial resolution unveil patterns of regional income segregation. Our proposed method of MERFs provides useful results with remarkably higher accuracy than direct estimates and the EBP-BC for most of out-of-sample domains. The following design-based simulation will




FIGURE 6 Domain-specific coefficients of variations for target variable *ictpc* for in- and out-of-sample domains [Colour figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com/doi/10.1111/rasc.12600)]

strengthen the reliability of results and enable an in-depth discussion of our methods for point and MSE-estimates.

5.3 | Evaluation using design-based simulation

The design-based simulation allows to directly juxtapose the performance of the proposed MERF-approach to existing SAE-methods for the estimation of area-level means based on empirical data. In this sense, the design-based simulation adds not only insights to the results from the model-based simulation in Section 4, but also evaluates results from the example in the previous Section 5.3. We focus on area-level mean-estimates of household income from work *inglabpc* in the Mexican state Nuevo León. As we use the same data with a different target variable, sample and census data properties are similar to the previous example with details provided in Table 4. Implementing the design-based simulation, we sample $T = 500$ independent samples from the fixed population of our census dataset. Each pseudo-survey-sample mirrors the characteristics of the original survey, as we keep the number of in-sample households similar to the original sample sizes and abstain from sampling out-of-sample municipalities. As a result, we use 500 equally structured pseudo-survey-samples with equal overall sample size. True values, are defined as the domain-level averages of household incomes from work in the original census.

We consider the same methods as in the model-based simulation in Section 4. Comparably to Section 5.2, we use the same working-model for the BHF, the EBP, the EBP-BC and P-SPLINES




FIGURE 7 Detailed comparison of point estimates for the domain-level average total household income. The dotted line separates sampled from non-sampled areas. In-sample areas are sorted by decreasing survey-sample size. [Colour figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com/doi/10.1111/rssc.12400)]

and assume it to be fixed throughout the design-based simulation. For the EBP-BC and the MERF, we keep the parameters as already discussed in Section 5.2.

The discussion of results, starts with an investigation into the performance of point estimates. Figure 8 reports the average RMSE of the area-level mean-estimates for Nuevo León in total and with details on the 21 in- and 30 out-of sample areas. The corresponding summary statistics for Figure 8 are given in Table B2 in Appendix B. Regarding the total of 51 areas, we observe no remarkable difference in the performance of the BHF and the EBP, whereas the EBP-BC has lower RMSE on average. P-SPLINES outperform the BHF, EBP and EBP-BC in mean and median terms of RMSE. The MERF point estimates indicate the lowest RMSEs among all areas resulting in an more than 22% improvement compared to the BHF. Referring to the RMSE for in-sample areas, we see two different ways how the EBP-BC and the adaptive methods of P-SPLINES and MERF deal with high and unbalanced variation in our true values for certain areas, ranging from 475 to about 4004 pesos: overall, the MERF deals best in modelling the complex survey data and produces highly accurate estimates for the majority of in-sample areas. A closer look at the results, reveals however, that higher RMSE values due to overestimation mainly occur in two areas, both characterised by a very low level of income (622 and 544 pesos, respectively). A similar observation can be made for P-SPLINES, although the MERF appears to reproduce the predictive relations more efficiently. In contrast, we observe the in-sample behaviour of the EBP-BC, which clearly opposes its superior overall performance. The EBP-BC appears to balance extreme estimates by producing slightly worse estimates for each individual in-sample areas, than

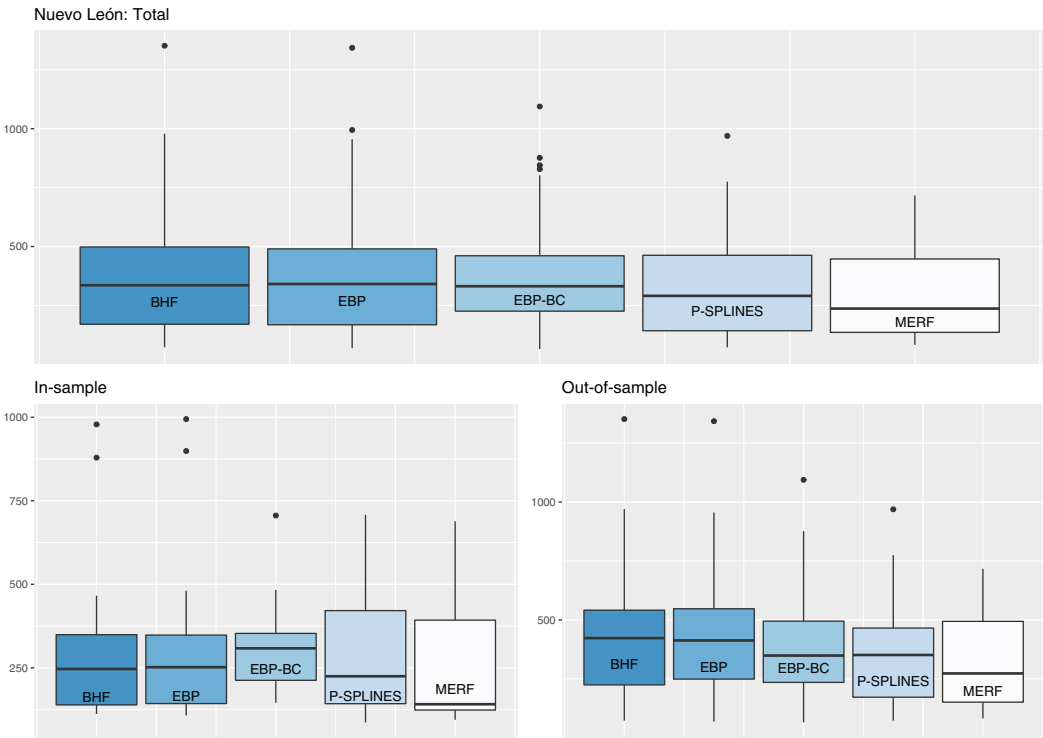


FIGURE 8 Performance of area-specific point estimates including details on in- and out-of-sample areas. Comparison of empirical root mean squared errors from the design-based simulation for target variable *inglabpc* [Colour figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com/doi/10.1111/rasc.12400)]

allowing for individually inferior estimates for specific ‘outlier’-areas. This behaviour is conceptually rooted in its data-driven transformation-approach. Nevertheless, this property enables the EBP-BC to identify a model, providing stable and precise estimates on the majority of areas, especially the 30 non-sampled areas. Given the data-scenario of Nuevo León, the performance on the out-of-sample areas, delineates each method’s quality and stability. In this case, the EBP-BC and the non-parametric approaches of P-SPLINES and MERF outperform the ‘traditional’ methods (BHF and EBP) in terms of lower RMSE. The median of RMSE of P-SPLINES aligns to the values of the EBP-BC, although the RMSEs of P-SPLINES are lower in means. One distinct advantage of the MERF is its adaptability and implicit model-selection that is rewarded in the presence of complex data-scenarios.

The findings from Figure 8, are strengthened by a discussion of mean and median values of RB and RRMSE in Table 5. Referring to all 51 areas, the RB of the data-driven method of EBP-BC, P-SPLINES, and the MERF is smaller in median values than the RB of BHF and the EBP. Respectively, the MERF shows comparatively low levels of median RB while mean values lie in the same range with competing methods. The obvious difference between mean and median values, indicates the previously discussed existence of inferior estimates for specific regions due to the empirical properties of our underlying data. For the 20 in-sample areas, P-SPLINES perform superior to competing methods regarding the median values of RB. The close relation between the mean and median values of RB for the EBP-BC highlight the mentioned balancing-property of the EBP-BC. For the majority of areas in the model-based simulation, that is, the 30 non-sampled areas, the EBP-BC, P-SPLINES as well as the MERF exhibit a comparatively low level of median

TABLE 5 Mean and Median of relative bias (RB) and relative root mean squared error (RRMSE) over in- and out-of-sample areas for point estimates

	Total		In-sample		Out-of-sample	
	Median	Mean	Median	Mean	Median	Mean
RB (%)						
BHF	14.200	17.800	5.070	11.200	21.800	22.400
EBP	14.700	17.900	5.310	11.700	22.600	22.300
EBP-BC	9.650	18.400	6.140	7.150	15.800	26.300
P-SPLINES	8.120	18.200	0.014	14.200	13.500	21.000
MERF	7.720	18.600	3.490	17.200	10.600	19.600
RRMSE (%)						
BHF	14.900	21.600	9.520	17.000	23.000	24.800
EBP	15.900	21.700	9.480	17.400	22.900	24.700
EBP-BC	14.000	23.900	12.900	15.200	16.100	29.900
P-SPLINES	11.600	22.800	7.360	20.000	15.000	24.700
MERF	9.430	21.200	6.130	19.900	12.500	22.100

Abbreviations: BHF, Battese et al.; EBP-BC, EBP under data-driven Box–Cox transformation; MERF, mixed effects random forests.

TABLE 6 Performance of mean squared-error (MSE)-estimator in design-based simulation: mean and median of relative bias (RB) and relative root mean squared error (RRMSE) over in- and out-of-sample areas

	Total		In-sample		Out-of-sample	
	Median	Mean	Median	Mean	Median	Mean
RB-RMSE (%)	−1.260	14.300	0.719	7.820	−9.460	18.800
RRMSE-RMSE (%)	48.100	55.900	41.400	47.700	50.900	61.700

RB. Especially the MERF captivates by the lowest values in median RRMSE compared to its competitors, while mean values for all cases are within a compatible range.

Finally, we focus on the performance of the proposed non-parametric MSE-bootstrap procedure. While, the model-based simulation in Section 4 indicates unbiasedness of the proposed bootstrap-scheme under all four scenarios, our results from the design-based simulation require a deeper discussion. We abstain from a discussion of result from our analytical approximation to the area-level MSE because it is limited to in-sample areas and was solely used to contextualise the quality of our proposed bootstrap-scheme from Section 3. Table 6 reports the results of RB-RMSE and the RRMSE-RMSE for the corresponding estimates. Figure B2 in the Appendix B visualises details from Table 6. First of all, the values of RRMSE-RMSE are comparable to the most complex scenario in the model-based scenario in Section 4. The RB-RMSE for the in-sample areas indicates unbiasedness in median terms and an acceptable overestimation regarding the mean RB-RMSE. For out-of-sample areas, we face a moderate underestimation regarding the median value and over-estimation according to mean values. Nevertheless, Figure B2 in Appendix B reveals, that the mixed signal between mean and median in Table 6 is explained by a balanced mix of under- and over-estimation. Overall, the expectations towards the MSE-bootstrap procedure, given the challenging conditions of this design-based simulation, are met. Especially, the results

for in-sample areas, combined with insights from the model-based simulation, indicate a solid and reliable performance of the proposed non-parametric bootstrap procedure. Although, the RB-RMSE for all 51 areas is driven by the results from out-of-sample areas, the median RB-RMSE is acceptable. Apparently, the MSE-estimates mirror the high variation in sample sizes paired with high and disproportional variation of high-income and low-income regions between the 21 in-sample and 30 out-of-sample areas. From an applied perspective, the MSE-estimates for out-of-sample areas are nevertheless practicable for the construction of confidence intervals, with a median coverage of 0.97.

6 | CONCLUDING REMARKS

In this paper, we explore the potential of tree-based machine learning methods for the estimation of SAE-means. In particular, we provide a solid framework easing the use of random forests for regression within the existing methodological framework of SAE. We highlight the potential of our approach to meet modern requirements of SAE, including the robustness of random forests against model-failure and the applicability for high-dimensional problems processing Big Data sources. The methodological part focusses on the MERF-procedure (Hajjem et al., 2014) and implicitly discusses a semi-parametric unit-level mixed model, treating LMM-based SAE-methods, such as the BHF and the EBP, as special cases. The model is fit by an algorithm resembling the EM-algorithm, allowing for flexibility in the specification to model fixed effects as well as random-effects. The proposed point estimator for area-level means is complemented by the non-parametric MSE-bootstrap scheme, building on the REB-bootstrap by Chambers and Chandra (2013) and the bias-corrected estimate for the residual variance by Mendez and Lohr (2011). We evaluate the performance of point- and MSE-estimates compared to 'traditional' SAE-methods by model- and design-based simulations and provide a distinctive SAE example using income data from the Mexican state Nuevo León in Section 5.2. The model-based simulation in Section 4 demonstrates the ability of point estimates to perform compatibly in classical scenarios and outperform 'traditional' methods in the existence non-linear interactions between covariates. The design-based simulation in Section 5.3 confirms the adequacy of MERFs for point estimates under searingly realistic conditions. The model- and design-based simulations indicate that the proposed approach is robust against distributional violations of normality for the random effects and for the unit-level error terms. Concerning our proposed MSE-bootstrap scheme, we conclude its reliability based on the performance in the model-based simulation in Section 4. Furthermore, we obtain reasonable support for the performance in the application in Section 5.2 and the following design-based simulation in Section 5.3.

We motivate three major dimensions for further research, including theoretical work, aspects of generalisations and advanced applications using Big Data covariates: from a theoretical perspective, further research is needed to investigate the construction and theoretical discussion of a partial-analytical MSE for area-level means. A conducive strategy is an extension based on our theoretical discussion in the online supplementary materials (S1). Additionally the deduction of recent theoretical results, such as conditions for the consistency of unit-level predictions (Scornet et al., 2015) or considerations of individual predictions intervals (Wager et al., 2014; Zhang et al., 2019), towards area-level indicators, bears potential. Alternatively, considerations concerning a fully non-parametric formulation of model (1) impose an interesting research direction. From a survey statistical perspective, our proposed method currently abstains from the use of survey weights which bears a risk if the assumption of non-informative sampling is violated. Nevertheless, there exist approaches incorporating weights into random forests

(Winham et al., 2013). The transfer of such ideas to the proposed method of MERFs is subject to ongoing research. Regarding additional generalisations of the proposed method, we aim to extend the use of MERFs towards the estimation of small area quantiles and other non-linear indicators, such as Gini-coefficients or Head Count Ratios. Furthermore, a generalisation towards binary or count data is possible and left to further research. The semi-parametric composite formulation of model (1) allows for $f(\cdot)$ to adapt any functional form regarding the estimation of the conditional mean of y_i given X_i and technically transfers to other machine learning methods, such as gradient-boosted trees or support vector machines. In terms of advanced applications, we propose the use of MERFs for complex random effect and covariance-structures in empirical problems to the SAE-research community. Equally interesting is the use of high-dimensional supplementary data, that is, Big Data covariates, for the estimation of area-level means, that can be directly handled by the proposed MERF-framework.

ACKNOWLEDGEMENTS

The authors are grateful to CONEVAL for providing the data used in empirical work. The views set out in this paper are those of the authors and do not reflect the official opinion of CONEVAL. The numerical results are not official estimates and are only produced for illustrating the methods. Additionally, the authors would like to thank the HPC Service of ZEDAT, Freie Universität Berlin, for computing time. Finally, the authors are indebted to the Joint Editor, Associate Editor and the referee for comments that significantly improved the paper. Open Access funding enabled and organized by Projekt DEAL.

DATA AVAILABILITY STATEMENT

Software and synthetic data are available on request from the authors.

ORCID

Timo Schmid  <https://orcid.org/0000-0002-7217-2501>

REFERENCES

- Alfons, A. & Templ, M. (2013) Estimation of social exclusion indicators from complex surveys: the R package Laeken. *Journal of Statistical Software*, 54(15), 1–25.
- Anderson, W., Guikema, S., Zaitchik, B. & Pan, W. (2014) Methods for estimating population density in data-limited areas: evaluating regression and tree-based models in Peru. *PLoS One*, 9(7), e100037.
- Athey, S., Tibshirani, J. & Wager, S. (2019) Generalized random forests. *The Annals of Statistics*, 47(2), 1148–1178.
- Bates, D., Mächler, M., Bolker, B. & Walker, S. (2015) Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48.
- 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.
- Biau, G. & Scornet, E. (2016) A random forest guided tour. *Test*, 25(2), 197–227.
- Bilton, P., Jones, G., Ganesh, S. & Haslett, S. (2017) Classification trees for poverty mapping. *Computational Statistics & Data Analysis*, 115, 53–66.
- Breiman, L. (1996) Bagging predictors. *Machine Learning*, 24(2), 123–140.
- Breiman, L. (2001) Random forests. *Machine Learning*, 45(1), 5–32.
- Breiman, L., Friedman, J., Stone, C.J. & Olshen, R.A. (1984) *Classification and regression trees*. Boca Raton: CRC Press.
- Capitaine, L., Genuer, R. & Thiébaud, R. (2021) Random forests for high-dimensional longitudinal data. *Statistical Methods in Medical Research*, 30(1), 166–184.
- Chambers, R. & Chandra, H. (2013) A random effect block bootstrap for clustered data. *Journal of Computational and Graphical Statistics*, 22(2), 452–470.

- Chambers, R. & Tzavidis, N. (2006) M-quantile models for small area estimation. *Biometrika*, 93, 255–268.
- Dagdoug, M., Goga, C. & Haziza, D. (2021) Model-assisted estimation through random forests in finite population sampling. *Journal of the American Statistical Association*, 1–18.
- Datta, G.S. & Lahiri, P. (2000) A unified measure of uncertainty of estimated best linear unbiased predictors in small area estimation problems. *Statistica Sinica*, 10(2), 613–627.
- De Moliner, A. & Goga, C. (2018) Sample-based estimation of mean electricity consumption curves for small domains. *Survey Methodology*, Statistics Canada, Catalogue No. 12-001-X, 44(2), 1–24.
- Diallo, M.S. & Rao, J.N.K. (2018) Small area estimation of complex parameters under unit-level models with skew-normal errors. *Scandinavian Journal of Statistics*, 45(4), 1092–1116.
- Efron, B. & Hastie, T. (2016) *Computer age statistical inference*. Cambridge: Cambridge University Press.
- 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.
- González-Manteiga, W., Lombardía, M.J., Molina, I., Morales, D. & Santamaría, L. (2008) Bootstrap mean squared error of a small-area EBLUP. *Journal of Statistical Computation and Simulation*, 78(5), 443–462.
- Graf, M., Marin, J.M. & Molina, I. (2019) A generalized mixed model for skewed distributions applied to small area estimation. *Test*, 28(2), 565–597.
- Greenwell, B.M. (2017) pdp: an R package for constructing partial dependence plots. *The R Journal*, 9(1), 421–436.
- Greenwell, B.M., Boehmke, B. & Gray, B. (2020) Variable importance plots—An introduction to the VIP package. *The R Journal*, 12(1), 343–366.
- Hajjem, A., Bellavance, F. & Larocque, D. (2011) Mixed effects regression trees for clustered data. *Statistics & Probability Letters*, 81(4), 451–459.
- Hajjem, A., Bellavance, F. & Larocque, D. (2014) Mixed-effects random forest for clustered data. *Journal of Statistical Computation and Simulation*, 84(6), 1313–1328.
- Hall, P. & Maiti, T. (2006) On parametric bootstrap methods for small area prediction. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(2), 221–238.
- Hastie, T., Tibshirani, R. & Friedman, J. (2009) *The elements of statistical learning: data mining, inference, and prediction*. New York: Springer Science & Business Media.
- Jiang, J. & Rao, J.S. (2020) Robust small area estimation: an overview. *Annual Review of Statistics and Its Application*, 7(1), 337–360.
- Kreutzmann, A.-K., Pannier, S., Rojas-Perilla, N., Schmid, T., Templ, M. & Tzavidis, N. (2019) The R package EMDI for estimating and mapping regionally disaggregated indicators. *Journal of Statistical Software*, 91(7), 1–33.
- Lambert, F. & Park, H. (2019) *Income inequality and government transfers in Mexico* (IMF Working Papers 148).
- Liaw, A. & Wiener, M. (2002) Classification and regression by randomforest. *R News*, 2(3), 18–22.
- Marchetti, S., Giusti, C., Pratesi, M., Salvati, N., Giannotti, F., Pedreschi, D. et al. (2015) Small area model-based estimators using big data sources. *Journal of Official Statistics*, 31(2), 263–281.
- Marchetti, S. & Tzavidis, N. (2021) Robust estimation of the theil index and the gini coefficient for small areas. *Journal of Official Statistics*, 37(4), 955–979.
- Marino, M.F., Ranalli, M.G., Salvati, N. & Alfo, M. (2019) Semi-parametric empirical best prediction for small area estimation of unemployment indicators. *Annals of Applied Statistics*, 13(2), 1166–1197.
- Marino, M.F., Tzavidis, N. & Alfo, M. (2018) Mixed hidden Markov quantile regression models for longitudinal data with possibly incomplete sequences. *Statistical Methods in Medical Research*, 27(7), 2231–2246.
- McConville, K.S. & Toth, D. (2019) Automated selection of post-strata using a model-assisted regression tree estimator. *Scandinavian Journal of Statistics*, 46(2), 389–413.
- Mendez, G. (2008) *Tree-based methods to model dependent data* (Unpublished doctoral dissertation). Arizona State University.
- Mendez, G. & Lohr, S. (2011) Estimating residual variance in random forest regression. *Computational Statistics & Data Analysis*, 55(11), 2937–2950.
- Molina, I. & Marhuenda, Y. (2015) sae: an R package for small area estimation. *The R Journal*, 7(1), 81–98.
- Molina, I. & Rao, J.N.K. (2010) Small area estimation of poverty indicators. *Canadian Journal of Statistics*, 38(3), 369–385.
- OECD. (2021) Income distribution. *OECD social and welfare statistics (database)*. Retrieved from: <https://doi.org/10.1787/data-00654-en> [Accessed 13th April 2021].

- Opsomer, J.D., Claeskens, G., Ranalli, M.G., Kauermann, G. & Breidt, F. (2008) Non-parametric small area estimation using penalized spline regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(1), 265–286.
- Prasad, N.G.N. & Rao, J.N.K. (1990) The estimation of the mean squared error of small-area estimators. *Journal of the American Statistical Association*, 85(409), 163–171.
- R Core Team. (2021) *R: a language and environment for statistical computing [Computer software manual]*. Vienna, Austria: R Foundation for Statistical Computing.
- Rao, J.N.K. & Molina, I. (2015) *Small area estimation. Wiley Series in Survey Methodology*, 2nd edition. Hoboken, NJ: Wiley.
- Rojas-Perilla, N., Pannier, S., Schmid, T. & Tzavidis, N. (2020) Data-driven transformations in small area estimation. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 183(1), 121–148.
- Schmid, T., Bruckschen, F., Salvati, N. & Zbiranski, T. (2017) Constructing sociodemographic indicators for national statistical institutes by using mobile phone data: estimating literacy rates in senegal. *Journal of the Royal Statistical Society: Series A*, 180(4), 1163–1190.
- Scornet, E., Biau, G. & Vert, J.-P. (2015) Consistency of random forests. *The Annals of Statistics*, 43(4), 1716–1741.
- Sela, R.J. & Simonoff, J.S. (2012) RE-EM trees: a data mining approach for longitudinal and clustered data. *Machine Learning*, 86(2), 169–207.
- Sexton, J. & Laake, P. (2009) Standard errors for bagged and random forest estimators. *Computational Statistics & Data Analysis*, 53(3), 801–811.
- Smits, J. & Permanyer, I. (2019) The subnational human development database. *Scientific Database*, 6, 190038.
- Sugasawa, S. & Kubokawa, T. (2017) Transforming response values in small area prediction. *Computational Statistics & Data Analysis*, 114, 47–60.
- Sugasawa, S. & Kubokawa, T. (2019) Adaptively transformed mixed-model prediction of general finite-population parameters. *Scandinavian Journal of Statistics*, 46(4), 1025–1046.
- Tzavidis, N., Marchetti, S. & Chambers, R. (2010) Robust estimation of small-area means and quantiles. *Australian & New Zealand Journal of Statistics*, 52(2), 167–186.
- Tzavidis, N., Zhang, L.-C., Luna, A., Schmid, T. & Rojas-Perilla, N. (2018) From start to finish: a framework for the production of small area official statistics. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 181(4), 927–979.
- Varian, H.R. (2014) Big data: new tricks for econometrics. *Journal of Economic Perspectives*, 28(2), 3–28.
- Wager, S. & Athey, S. (2018) Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523), 1228–1242.
- Wager, S., Hastie, T. & Efron, B. (2014) Confidence intervals for random forests: the jackknife and the infinitesimal jackknife. *The Journal of Machine Learning Research*, 15(1), 1625–1651.
- Winham, S.J., Freimuth, R.R. & Biernacka, J.M. (2013) A weighted random forests approach to improve predictive performance. *Statistical Analysis and Data Mining*, 6(6), 496–505.
- Wood, S. (2017) *Generalized additive models: An introduction with R*, 2nd edition. Boca Raton: Chapman & Hall/CRC Press.
- Wu, H. & Zhang, J.-T. (2006) *Nonparametric regression methods for longitudinal data analysis: mixed-effects modeling approaches*. Hoboken, NJ: John Wiley & Sons.
- Zhang, H., Zimmerman, J., Nettleton, D. & Nordman, D.J. (2019) Random forest prediction intervals. *The American Statistician*, 74(4), 392–406.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.

How to cite this article: 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. Available from: <https://doi.org/10.1111/rssc.12600>

APPENDIX A

After convergence of the algorithm introduced in Section 2.2, we obtain an optimal non-parametric estimator for $\hat{f}(\cdot)$. In the following, we facilitate the notation and refer to $\hat{f}^{\text{OOB}}(\cdot)$ simply as $\hat{f}(\cdot)$. The best predictor for the random effects for known parameters H_i and R_i must maximise the GLL criterion:

$$\text{GLL}(f, v_i | y) = \sum_{i=1}^D \{ [y_i - f(X_i) - Z_i v_i]' R_i^{-1} [y_i - f(X_i) - Z_i v_i] + v_i' H_i^{-1} v_i + \log |H| + \log |R_i| \}.$$

Finding a maximum for GLL is equivalent to the problem of finding a maximiser for v in the first term of the summation of the proposed GLL-criterion:

$$[y - \hat{f}(X_i) - Z_i v_i]' R_i^{-1} [y_i - \hat{f}(X_i) - Z_i v_i] + v_i' H_i^{-1} v_i.$$

Reshaping leads to:

$$\begin{aligned} & [y_i - \hat{f}(X_i) - Z_i v_i]' R_i^{-1} [y_i - \hat{f}(X_i) - Z_i v_i] + v_i' H_i^{-1} v_i \\ & = y_i' R_i^{-1} y_i - y_i' R_i^{-1} \hat{f}(X_i) - y_i' R_i^{-1} Z_i v_i - \hat{f}(X_i)' R_i^{-1} y_i + \hat{f}(X_i)' R_i^{-1} \hat{f}(X_i) \\ & \quad + \hat{f}(X_i)' R_i^{-1} Z_i v_i - (Z_i v_i)' R_i^{-1} y_i + (Z_i v_i)' R_i^{-1} \hat{f}(X_i) + (Z_i v_i)' R_i^{-1} (Z_i v_i) + v_i' H_i^{-1} v_i. \end{aligned}$$

Now we derive the expression with respect to v and set it to 0 in order to find the maximiser

$$\begin{aligned} & -y_i' R_i^{-1} Z_i + \hat{f}(X_i)' R_i^{-1} Z_i - Z_i' R_i^{-1} y_i + \\ & Z_i' R_i^{-1} \hat{f}(X_i) + Z_i' R_i^{-1} Z_i v_i + (Z_i v_i)' R_i^{-1} y_i + Z_i + 2H_i^{-1} v_i = 0 \quad \Leftrightarrow \\ & -2y_i R_i^{-1} Z_i + 2Z_i' R_i^{-1} \hat{f}(X_i) + 2Z_i' R_i^{-1} Z_i v_i + 2H_i^{-1} v_i = 0 \quad \Leftrightarrow \\ & -y_i R_i^{-1} Z_i + Z_i' R_i^{-1} \hat{f}(X_i) + Z_i' R_i^{-1} Z_i v_i + H_i^{-1} v_i = 0 \quad \Leftrightarrow \\ & \quad \quad \quad y_i R_i^{-1} Z_i - Z_i' R_i^{-1} \hat{f}(X_i) = Z_i' R_i^{-1} Z_i v_i + H_i^{-1} v_i \quad \Leftrightarrow \\ & \quad \quad \quad y_i R_i^{-1} Z_i - Z_i' R_i^{-1} \hat{f}(X_i) = (Z_i' R_i^{-1} Z_i + H_i^{-1}) v_i \quad \Leftrightarrow \\ & (Z_i' R_i^{-1} Z_i + H_i^{-1})^{-1} (Z_i' R_i^{-1} (y_i - \hat{f}(X_i))) = v_i \\ & v_i = (Z_i' R_i^{-1} Z_i + H_i^{-1})^{-1} (Z_i' R_i^{-1} (y_i - \hat{f}(X_i))) \\ & = (Z_i' R_i^{-1} Z_i + H_i^{-1})^{-1} V_i V_i^{-1} (Z_i' R_i^{-1} (y_i - \hat{f}(X_i))) \\ & = (Z_i' R_i^{-1} Z_i + H_i^{-1})^{-1} Z_i' R_i^{-1} (R_i + Z_i H_i Z_i') V_i^{-1} (y_i - \hat{f}(X_i)) \\ & = (Z_i' R_i^{-1} Z_i + H_i^{-1})^{-1} (Z_i' + Z_i' R_i^{-1} Z_i H_i Z_i') V_i^{-1} (y_i - \hat{f}(X_i)) \\ & = (Z_i' R_i^{-1} Z_i + H_i^{-1})^{-1} (Z_i' R_i^{-1} Z_i + H_i^{-1}) (H_i Z_i' V_i^{-1} (y_i - \hat{f}(X_i))) \\ & = (H_i Z_i' V_i^{-1} (y_i - \hat{f}(X_i))). \end{aligned}$$

The solution of the maximisation problem is given by $\hat{v}_i^* = H_i Z_i' V_i^{-1} (y_i - \hat{f}(X_i))$. Note, for $\hat{f}(X_i) = X_i \hat{\beta}$, the optimality solution resembles the BLUP.

APPENDIX B ADDITIONAL SIMULATION RESULTS AND MODEL-DIAGNOSTICS




FIGURE B1 Details on the performance of the proposed bootstrap mean squared error (MSE)-estimator and the analytic approximation in the model-based simulation: boxplots of the area-specific relative bias-root mean squared errors (RB-RMSEs) averaged over simulation runs [Colour figure can be viewed at wileyonlinelibrary.com]

TABLE B1 Explanation of most influential variables according to the random forest model in the application of Section 5.2

Variable name	Explanation
ictpc	Total household income per capita
escol_rel_hog	Average relative amount of schooling standardised by age and sex of household members
bienes	Availability of goods in the household
jaesc	Average years of schooling of persons in the household
jnived	Formal education of the head of the household
actcom	Assets in the household
pcpering	Percentage of income earners in the household
jexp	Years of working experience of the head of the household
pcpering_2	Number of income earners in the household by household size
pcocup	Percentage of people employed in the household
jtocup	Occupation type

TABLE B2 Performance of point estimates in design-based simulation: summary statistics of empirical root mean squared error for area-level mean-estimates

Areas	Method	Min	First Qu.	Median	Mean	Third Qu.	Max
Total	BHF	72.58	170.63	336.06	386.40	498.13	1351.82
	EBP	68.83	168.13	341.10	387.91	490.14	1342.92
	EBP-BC	65.22	225.84	331.64	376.49	460.86	1094.08
	P-SPLINES	72.70	142.91	290.84	337.11	462.94	969.15
	MERF	82.58	136.01	236.86	298.20	447.21	716.72
In-sample	BHF	111.93	139.13	246.62	301.90	349.26	978.49
	EBP	107.41	143.07	251.95	308.47	348.00	994.45
	EBP-BC	145.48	212.64	308.43	314.08	353.02	705.81
	P-SPLINES	86.69	142.70	224.82	285.37	421.29	707.71
	MERF	94.56	123.72	141.16	264.27	392.73	688.74
Out-of-sample	BHF	72.58	224.35	422.90	445.55	541.31	1351.82
	EBP	68.83	248.89	412.91	443.52	547.12	1342.92
	EBP-BC	65.22	234.98	348.83	420.17	494.83	1094.08
	P-SPLINES	72.70	172.29	351.17	373.34	465.36	969.15
	MERF	82.58	151.24	272.97	321.95	493.85	716.72

Abbreviations: BHF, Battese et al.; EBP-BC, EBP under data-driven Box-Cox transformation; MERF, mixed effects random forests.

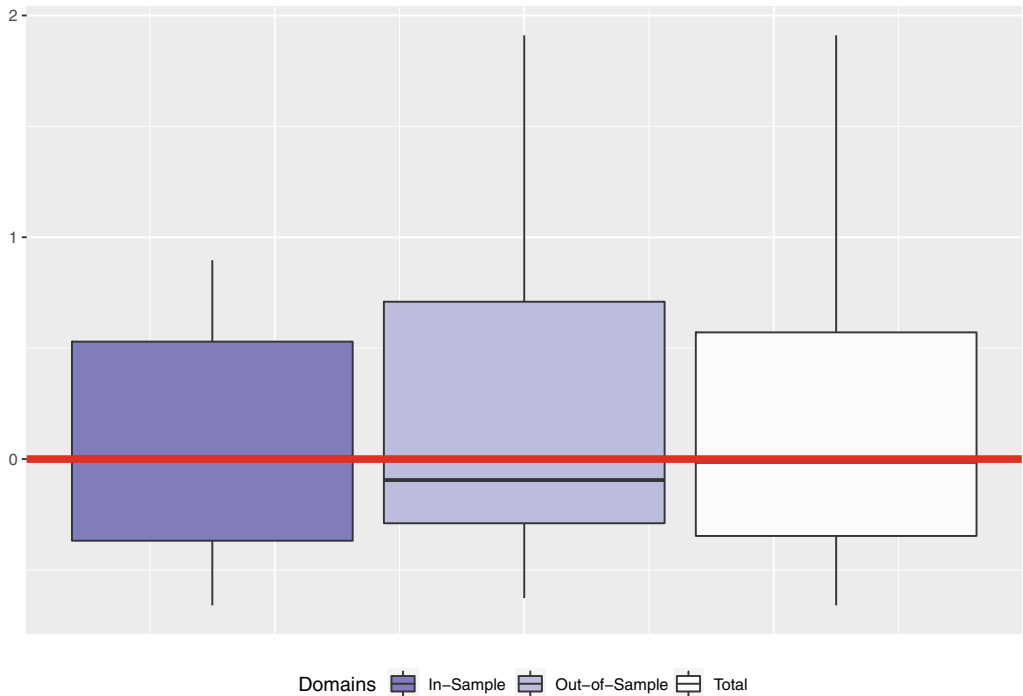


FIGURE B2 Details on the performance of the proposed mean squared error -estimator in the design-based simulation: boxplots of the area-specific relative bias-root mean squared errors averaged over simulation runs including details on in- and out-of-sample areas [Colour figure can be viewed at wileyonlinelibrary.com]