

On the performance of small-area estimators: fixed vs. random area parameters*

Alex Costa¹, Albert Satorra and Eva Ventura²

Abstract

Most methods for small-area estimation are based on composite estimators derived from design- or model-based methods. A composite estimator is a linear combination of a direct and an indirect estimator with weights that usually depend on unknown parameters which need to be estimated. Although model-based small-area estimators are usually based on random-effects models, the assumption of fixed effects is at face value more appropriate. Model-based estimators are justified by the assumption of random area effects; in practice, however, areas can not be substituted for one another in a random manner (we say, they are not interchangeable). In the present paper we empirically assess the quality of several small-area estimators in the setting in which the area effects are treated as fixed. We consider two settings: one that draws samples from a theoretical population, and another that draws samples from an empirical population of a labour force register maintained by the National Institute of Social Security (NISS) of Catalonia. We distinguish two types of composite estimators: *a*) those that use weights that involve area specific estimates of bias and variance; and, *b*) those that use weights that involve a common variance and a common squared bias estimate for all the areas. We assess their precision and discuss alternatives to optimizing composite estimation in applications.

MSC: 62G10,62J02

Keywords: Small area estimation, composite estimator, Monte Carlo study, random effect model, BLUP, empirical BLUP

1 Introduction

Sample surveys are often used to estimate quantities related not only to the total population but also to a variety of small-area domains. Small-area estimation is concer-

* This research was supported by IDESCAT (Statistical Institute of Catalonia). The views expressed in this paper are those of the authors and do not necessarily represent the policies of IDESCAT. Detailed and very helpful comments by Nicholas T. Longford on a previous version of this paper are acknowledged. Constructive and detail reviewers' comments are also appreciated. Partial support by the Spanish Ministry of Science and Technology grant SEJ2006-13537 to A. Satorra and E. Ventura is also acknowledged.

¹ Idescat (Statistical Institute of Catalonia).

² Department of Economics and Business, Universitat Pompeu Fabra and Barcelona GSE.

Received: February 2008

Accepted: December 2008

ned with estimating population quantities associated with a partition of the domain (population) into subdomains (small areas or districts) $j = 1, \dots, J$. Nowadays there is a large body of methodology for small-area estimation; see, e.g., Platek, Rao, Särndal and Singh (1987), Isaki (1990), Ghosh and Rao (1994), Singh, Gambino and Mantel (1994), and Rao (2003).

Large-scale (national) surveys are usually designed to yield estimates of a small number of key national population quantities (means, proportions and the like) that have sufficient precision, without having to adopt any assumptions other than the sampling design. Insisting on a large sample for each district is not realistic, especially when there are many districts, and several of them form a very small fraction of the population.

When estimating a domain quantity, we refer to a *direct estimator* if it is based only on the domain-specific sample. A domain (area) is regarded as small if the direct estimate for the area does not have adequate precision. For a small area one could use *indirect estimators* that borrow strength from values of the variable of interest from related areas and/or time periods. An implicit or explicit model is used to link the different areas and/or time periods, often through the use of auxiliary information such as a census count or some administrative records. An initial classification of small-area estimation divides the methods into *design-based* and *model-based*.

Design-based methods are based solely on the sampling design and do not make use of distributional (model) assumptions about the observed variables. Sampling variation, that is, variation across hypothetical replications of drawing a sample, arises only due to the variation of the specific units that are selected into the sample, and not due to variation of the population characteristics of interest (such as the small-area means) which are considered fixed because they are constant across replications. In contrast, model-based methods assume stochastic models governing the population values that are the target of the estimation process. Models are used to mediate the process of borrowing strength across the districts (small areas). That is, inference about a district that is represented in the sample by very few observations is supported by the information in the other districts' subsamples. This is most effective when the districts are very similar. Similarity can be enhanced by adjustment for other variables, opening up the potential of regression models.

Borrowing strength, as defined originally by Efron and Morris (1973), is based on the assumption of random effects. In the simplest setting with no covariates, the deviations of the district-level means θ_j from their national mean θ are assumed to be a random sample from a centred distribution with a finite variance, such as $\mathcal{N}(0, \sigma_u^2)$.

Model-based methods for small-area estimation associate the districts with random effects. In applications, however, the districts have their names (labels), and the target quantities θ_j could in principle be established by enumeration. In an hypothetical replication of the survey, the same districts, with the same subpopulations and the same values of θ_j would be involved. Therefore, it is natural to associate the districts with fixed effects. Longford (2007) argues that the assumption of fixed or random effect has a profound effect on standard errors of model-based small-area estimators. In the present paper we consider both design- and model-based estimators, and assess their accuracy

in the case of the fixed-effect assumption. Accuracy refers not to average MSE across areas, but to MSE for the particular (fixed) areas. This departs from previous studies in which accuracy was assessed by averaging MSE across areas (see, e.g., Costa, Satorra and Ventura (2003), and Santamaría, Morales and Molina (2004)).

In the model-based approach, the best linear unbiased predictor (BLUP) of the parameter of interest (the small-area parameter), is a linear combination of a direct and a synthetic estimator with weights that depend on two parameters that are usually unknown: the within- and between-area variances (possibly after controlling for other variables, regressors). Since both parameters are unknown quantities, these two variances have to be estimated, giving rise to the empirical BLUP (EBLUP). This estimation can distort the optimality of the EBLUP. In sections 3 and 4 we assess the consequences on accuracy of the substitution of model parameters by estimated values.

The purpose of the paper is to compare the performance of design- vs. model-based small-area estimators, with a focus on a specific (fixed) set of small areas. Monte Carlo methods are used for this investigation.

Two population frames will be considered in the Monte Carlo study: *a*) a theoretical population with varying distribution and sample size; *b*) an empirical population of labour statistics from the affiliation of firms in the NISS (National Institute of Social Security) registers. The choice of the NISS is motivated by current work at IDESCAT (Statistics Bureau of Catalonia).

The plan of the paper is as follows. Section 2 develops the notation and general context of small-area estimation, focusing on the distinction between design-based and model-based methods. Sections 3 and 4 describe the Monte Carlo studies using the theoretical and the empirical population, respectively. Section 5 concludes with a discussion of the results and the avenues for further research.

2 Approaches for small-area estimation

We consider a population stratified into J (small-area) domains (strata), $j = 1, 2, \dots, J$, and we seek to estimate the stratum parameters θ_j as well as an overall population parameter θ . A direct estimator of θ_j uses sample data only from area j . An indirect or synthetic estimator of θ_j uses data also from outside area j . We suppose that there is a direct estimator $\hat{\theta}_{dj}$ of θ_j and that it is unbiased (but may have large variance), and a synthetic estimator $\hat{\theta}_{sj}$ that has small variance but may be biased for θ_j .

Two perspectives motivate the different small-area estimators. The first assumes that the θ_j are fixed values and that there is sampling variation only within each stratum. In the second, in addition to the random variation within strata, there is also random variation of the θ_j , that are supposed to be realizations from a specific sampling distribution. We now describe these two approaches, design-based (fixed θ_j) and model-based (random θ_j), respectively.

2.1 Fixed-area perspective

Following Rao (2003, Section 4.3), a natural way to balance the potential bias of a synthetic estimator $\hat{\theta}_{sj}$ of θ_j against the instability of a direct estimator $\hat{\theta}_{dj}$ of the same parameter is to take the composite estimator (weighted average)

$$\hat{\theta}_{cj}(\pi_j) = (1 - \pi_j)\hat{\theta}_{dj} + \pi_j\hat{\theta}_{sj}, \quad (1)$$

a function of the weight $0 \leq \pi_j \leq 1$. This estimator has a mean square error (MSE) given by (Rao, 2003, formula (4.3.2)):¹

$$\begin{aligned} \text{MSE}(\hat{\theta}_{cj}, \theta_j) &= (1 - \pi_j)^2 \text{MSE}(\hat{\theta}_{dj}, \theta_j) + \pi_j^2 \text{MSE}(\hat{\theta}_{sj}, \theta_j) \\ &\quad + 2\pi_j(1 - \pi_j) \text{E} \left\{ (\hat{\theta}_{dj} - \theta_j)(\hat{\theta}_{sj} - \theta_j) \right\} \end{aligned} \quad (2)$$

where $\text{MSE}(\hat{\delta}, \delta)$ denotes the MSE of an estimator $\hat{\delta}$ with respect to the target δ . The expectation in the last term of (2) is taken with respect to the design-based sampling variation. In most applications, $\hat{\theta}_{dj}$ and $\hat{\theta}_{sj}$ are uncorrelated, so this last term vanishes. This is assumed throughout. Denote $\tilde{\theta}_{cj} = \hat{\theta}_{cj}(\tilde{\pi}_j)$.

The weight that minimizes the MSE of $\tilde{\theta}_{cj}$ is approximately (see Rao (2003, formula (4.3.3))

$$\tilde{\pi}_j = \frac{\text{MSE}(\hat{\theta}_{dj}, \theta_j)}{\text{MSE}(\hat{\theta}_{dj}, \theta_j) + \text{MSE}(\hat{\theta}_{sj}, \theta_j)} \quad (3)$$

in which case the (minimum) MSE is

$$\text{MSE}(\tilde{\theta}_{cj}, \theta_j) = \tilde{\pi}_j^2 \text{MSE}(\hat{\theta}_{sj}, \theta_j); \quad (4)$$

and

$$\text{MSE}(\tilde{\theta}_{cj}, \theta_j) = (1 - \tilde{\pi}_j)^2 \text{MSE}(\hat{\theta}_{dj}, \theta_j); \quad (5)$$

so, the (optimal) composite estimator $\tilde{\theta}_{cj}$ is superior to both the synthetic estimator $\hat{\theta}_{sj}$, since $\tilde{\pi}_j < 1$, and the direct estimator $\hat{\theta}_{dj}$, since $\tilde{\pi}_j > 0$ and $\tilde{\theta}_{cj}(0) = \hat{\theta}_{dj}$. If there was covariation among the synthetic and the direct estimator, $\text{cov}(\hat{\theta}_{dj}, \hat{\theta}_{sj})$ would be subtracted once in the numerator and twice in the denominator.

The expression (2) (with the covariance term ignored) will be used in sections 3 and 4 to compute the exact MSE of various composite estimators arising in a Monte Carlo study. The exact values of the MSE can be computed since in Monte Carlo studies we

¹ As in Rao (2003, Section 4.3), and throughout this section, “Var”, “MSE” and “E” should carry a subscript p that refers to variation with respect to the sampling design; this subscript has been suppressed for notational convenience.

know the population values of the parameters. In applications, the MSE will have to be estimated and several estimates are available. Longford (2007) discusses issues arising in the estimation of the MSE in the case of fixed area effects.

When the direct and synthetic estimators are unbiased for θ_j and θ respectively, and the variance of $\hat{\theta}_{sj}$ is small relative to the variance of the direct estimator, we have

$$\tilde{\pi}_j = \frac{\text{var}(\hat{\theta}_{dj})}{\text{var}(\hat{\theta}_{dj}) + (\theta - \theta_j)^2} \quad (6)$$

If $\hat{\theta}_{dj}$ is the sample mean, then $\text{var}(\hat{\theta}_{dj}) = \sigma_{j\epsilon}^2/n_j$, where $\sigma_{j\epsilon}^2$ is a within-domain variance and n_j is the sample size of the j th domain.² Then (6) becomes

$$\tilde{\pi}_j = \frac{\sigma_{j\epsilon}^2/n_j}{\sigma_{j\epsilon}^2/n_j + (\theta - \theta_j)^2} \quad (7)$$

For a synthetic estimator (unbiased for θ) whose variance is small compared with the variance of the direct estimator (unbiased for θ_j), we have

$$E(\hat{\theta}_{sj} - \hat{\theta}_{dj})^2 \approx (\theta - \theta_j)^2 + \text{var}(\hat{\theta}_{dj}) \quad (8)$$

and $\tilde{\pi}_j \approx \text{var}(\hat{\theta}_{dj})/(\hat{\theta}_{sj} - \hat{\theta}_{dj})^2$, suggesting the weight

$$\hat{\pi}_j^\dagger = \frac{\widehat{\text{var}}(\hat{\theta}_{dj})}{(\hat{\theta}_{sj} - \hat{\theta}_{dj})^2},$$

where $\widehat{\text{var}}(\hat{\theta}_{dj})$ is an unbiased estimator of $\text{var}(\hat{\theta}_{dj})$. This estimator is very unstable and it could even fall outside the interval $[0, 1]$. In the Monte Carlo study in sections 3 and 4 we use instead the weight

$$\hat{\pi}_j^* = \frac{\widehat{\text{var}}(\hat{\theta}_{dj})}{(\hat{\theta}_{sj} - \hat{\theta}_{dj})^2 + \widehat{\text{var}}(\hat{\theta}_{dj})}, \quad (9)$$

which satisfies the condition $0 \leq \hat{\pi}_j^* \leq 1$.

The composite small-area estimators can be based on the assumption of homogeneity of the within-area variances, in which case they will use a common estimate of this variance, such as the estimator of (14) defined in the section below, or they may contemplate heteroscedasticity, in which case they may use an area specific estimate such as (16) (see section below).

² We could contemplate homoscedasticity and replace $\sigma_{j\epsilon}^2$ by σ_ϵ^2 .

The optimal composite estimator that uses the weight in (7) is not feasible in practice because the bias term $(\theta_j - \theta)^2$ and the variance $\sigma_{j\epsilon}^2$ are unknown quantities that need to be estimated. We shall see that several alternatives to the estimation of the within-area variance do not induce much difference among estimators; in contrast, alternatives to the estimation of the squared area-bias term will lead to fundamental differences among estimators.

2.2 Random-area perspective

Alternative small-area estimators are based on models. Suppose

$$y_{ji} = X_{ji}\beta + Z_{ji}\gamma_j + \epsilon_{ji} \quad (10)$$

where $i = 1, 2, \dots, n_j$ and $j = 1, 2, \dots, J$, i and j denoting primary and secondary level units, observations and areas, respectively. X_{ji} and Z_{ji} are vectors of attributes of observation i of area j , β is a vector of regression coefficients and γ_j is a vector of random area effects, independent of ϵ_{ji} , and usually both normally distributed with respective variances σ_u^2 and σ_ϵ^2 (variance matrix Σ_u , instead of σ_u^2 , when Z_{ji} is a vector).

Since the issues we want to investigate arise already on the simplest of the random-area models, the one that has no covariates, our research will be made more transparent by using the simplest version of the model in (10), with Z_{ji} set to the indicator of area j and $X_{ji} = 1$ is empty (there are no covariates); that is,

$$y_{ji} = \mu + u_j + \epsilon_{ji} \quad (11)$$

Variables u_j and ϵ_{ij} are centred random variables with respective variances σ_u^2 (variance “between”) and σ_ϵ^2 (variance “within”).

Let $y_{j.} = n_j^{-1} \sum_i y_{ij}$ and $y_{..} = n^{-1} \sum_i \sum_j y_{ij}$ be the respective direct and synthetic estimators of $\theta_j = \mu + u_j$, where $n = \sum_j n_j$ is the overall sample size. Since $\text{var}(y_{j.}) = \sigma_u^2 + (\sigma_\epsilon^2/n_j)$ and $\text{cov}(y_{j.}, u_j) = \sigma_u^2$, the best unbiased linear predictor (BLUP³) of θ_j given $y_{j.}$ is (see, e.g., Neudecker and Satorra, 2003)

$$\text{BLUP}(\theta_j | y_{j.}) = \mu + \frac{\sigma_u^2}{\sigma_u^2 + \sigma_\epsilon^2/n_j} (y_{j.} - \mu) = (1 - \omega_j) y_{j.} + \omega_j \mu \quad (12)$$

where

$$\omega_j = 1 - \frac{\sigma_u^2}{\sigma_u^2 + \sigma_\epsilon^2/n_j} = \frac{\sigma_\epsilon^2/n_j}{\sigma_u^2 + \sigma_\epsilon^2/n_j} = \frac{1}{1 + n_j \gamma}$$

³ A common notation is also BLP, but since BLP is unbiased in the predictive sense, i.e. $E\{\text{BLP}(\theta_j) - \theta_j\} = 0$, the terminology of ‘best linear unbiased predictor’ (BLUP) will be used.

and $\gamma = \sigma_u^2/\sigma_\epsilon^2$. We used $E(y_{j.}) = \mu$. The empirical BLUP (EBLUP) is

$$\hat{\theta}_{cj}(\hat{\omega}) = \text{EBLUP}(\theta_j | y_{j.}) = (1 - \hat{\omega}_j)y_{j.} + \hat{\omega}_j y_{..}$$

where $y_{..}$ (the overall mean) is used as an estimator of μ and

$$\hat{\omega}_j = \frac{\hat{\sigma}_\epsilon^2/n_j}{\hat{\sigma}_u^2 + \hat{\sigma}_\epsilon^2/n_j} = \frac{1}{1 + n_j \hat{\gamma}} \quad (13)$$

as the estimator of ω_j , with $\hat{\gamma} = \hat{\sigma}_u^2/\hat{\sigma}_\epsilon^2$. Here

$$\hat{\sigma}_\epsilon^2 = \frac{1}{n - J} \sum_{j=1}^J \sum_{i=1}^{n_j} (y_{ij} - y_{j.})^2 \quad (14)$$

and

$$\hat{\sigma}_u^2 = \frac{1}{J - 1} \sum_{j=1}^J (y_{j.} - y_{..})^2 \quad (15)$$

are moment-matching estimators of the variances. The estimator $\hat{\sigma}_\epsilon^2$ could be alternatively written as a weighted mean, i.e.

$$\hat{\sigma}_\epsilon^2 = \sum_{j=1}^J ((n_j - 1)/(N - J)) \hat{\sigma}_{\epsilon j}^2,$$

of the within-area variance estimates

$$\hat{\sigma}_{j\epsilon}^2 = \frac{1}{n_j - 1} \sum_{i=1}^{n_j} (y_{ij} - y_{j.})^2. \quad (16)$$

As an alternative to these estimators we could use maximum likelihood (ML) estimation of the mixed regression model. For the unbalanced case, this provides alternative EBLUP estimators. These will be evaluated in the Monte Carlo study of sections 3 and 4.

One could also question the quality of these EBLUP estimators when the model (11) deviates from the standard assumptions, such as normality of the within- and between-area distributions, or both, or when there is variation among the within-area variances while equality is assumed.

Note that all composite estimators we consider are ‘‘borrowing strength’’ estimators, with the distinction that the ones based on the ‘‘random effect’’ perspective use ‘‘average’’ type of estimates for the bias, while the ones based on the ‘‘fixed effect’’ assumption use area specific estimate of the bias. In the Monte Carlo study below, we will see that there is a sharp difference on performance for these two type of estimators.

The Monte Carlo study of Section 3 contemplates normal and highly skewed distributions, both for the first- and second-level distributions. Non-normality of the distribution within each area, and heteroscedasticity of the within-area variances, is present in the Monte Carlo study of Section 4 involving an empirical population.

3 Monte Carlo study: theoretical population

In the simplest set-up, data is generated from a two-level model in which the domain parameters θ_j are realizations of $\theta_j \sim N(\mu = \theta, \sigma_u^2 = 3)$ and the observations y_{ji} (subject i in area j) are realizations of $y_{ji} \sim N(\mu = \theta_j, \sigma_\epsilon^2 = 6)$. The number of small areas is 40. In some simulations the within-area sample sizes are equal to $n_j = 10$, while in other simulations n_j ranges from 6 to 40.

Next we list the estimators considered in the Monte Carlo study. The direct and synthetic estimators are respectively the sample mean of area j and the overall sample mean $\hat{\theta}$. The composite estimators can be classified according to whether or not the weights are known (*theoretical*) or estimated (*empirical*), and according to whether the estimator of the squared bias term $(\theta_j - \theta)^2$ is *area specific* (weights will be denoted by π_j) or *averaged* across the areas (weights denoted as ω_j). Except for the direct estimator, denoted by D, all estimators considered are composite estimators whose weights are specified as follows:

DESIGN-BASED ESTIMATORS

Theoretical composite: TC1

$$\tilde{\pi}_j = \frac{\sigma_{j\epsilon}^2/n_j}{(\theta_j - \theta)^2 + \sigma_{j\epsilon}^2/n_j}$$

Empirical composite: CA

$$\hat{\pi}_j^* = \frac{\hat{\sigma}_{j\epsilon}^2/n_j}{(\hat{\theta}_j - \hat{\theta})^2 + \hat{\sigma}_{j\epsilon}^2/n_j}$$

Note that TC1 and CA use area-specific values for the within-area variance $\sigma_{j\epsilon}^2$ (they allow for heteroscedasticity of this variance across areas).

MODEL-BASED ESTIMATORS

Theoretical composite: TC2

$$\tilde{\omega}_j = \frac{\sigma_\epsilon^2/n_j}{\sigma_u^2 + \sigma_\epsilon^2/n_j},$$

with population (true) values for the variances within σ_ϵ^2 and between σ_u^2 .

Empirical composites: C

C is the composite estimator with $\hat{\omega}_j$ defined in (13), (14) and (15).

ML estimator: CML

Uses the estimator (12) with the population values μ , σ_ϵ^2 and σ_u^2 substituted by estimates obtained by fitting the model in (10) by ML.

3.1 Monte Carlo study: θ_j random

We first generate the area-level quantities θ_j as random draws from an assumed distribution $N(\theta, \sigma_u^2)$, independently across replications. The areas cannot be distinguished by any features (they are exchangeable), and so their MSEs are the same for all the areas, for each estimator. As should be expected, the results summarized in Figure 1 indicate that the MSEs for the different methods are highly correlated. Within a method, the empirical MSE's are not constant because the number of replications is finite (it is 3000).




Figure 1: Root MSE (RMSE) for areas $j = 1, \dots, 40$ when each mean θ_j is random across replications. Within- and between-area distributions are normal (number of replications is 3000, sample size in each area is 10). The mean and the median of the RMSE across areas are shown in the legend.

The theoretical design-based estimator TC1 is far more efficient than the others since it uses more information about the true values of the within-area variance and area-bias in each replication. The within-area variance is constant across replications, but this is not the case for the area-bias. The theoretical estimator TC2 that uses variance and bias parameters common across the areas performs similarly as the C and CML estimators (the last two estimators are equivalent, given that the n_j are equal across areas), which are the next in performance. The feasible design-based estimator CA performs poorly. Finally, the direct estimator has the poorest performance.

The gain of TC1 over TC2 can be explained by the fact that TC1 uses information about the squared area-bias $(\theta_j - \theta)^2$, which varies across replications, while TC2 uses only information about the true value of its expectation, the between-area variance (model parameter σ_u^2). Replacement of the parameters by their estimates in the model-based methods does not reduce this efficiency substantially; indeed, the RMSEs of TC2, C and CML are nearly indistinguishable in Figure 1. In contrast, the design-based estimator CA, which is based on substituting an estimate for the true value of the area-bias, incurs a severe loss of efficiency when compared with the theoretical estimator TC1.

3.2 Monte Carlo study: θ_j fixed

Now we assume that the θ_j are fixed across replications, in accordance with the empirical set-up in which the *eccentricity* of an area, i.e. the deviation of the area from the overall mean, is an (unknown) but fixed quantity that remains constant across replications.

Figure 2 reports the empirical root-MSE (RMSE) across replications for each area and for the different estimates. TC1 and TC2 are not feasible in practice since they use true values of population parameters that are not available in a typical application. However, the performance of TC1 and TC2 will shed light on the nature of the accuracy of the alternative estimators.

We see that the theoretical composite estimator TC1 that uses area-specific bias performs better than the theoretical composite TC2 that uses a single parameter (the variance between) to account for the squared bias averaged across the areas. In fact, TC2 performs poorly in all the areas, and as the worst estimator for areas with an extreme value of eccentricity (on the far right-hand side of the x-axis), that is, the areas for which the mean deviates highly from the overall area mean. To understand the variation of RMSE across the areas, these have been ordered according to their absolute deviation $|\theta_j - \theta|$, so that the extreme areas are located at the right-hand side. The lengths of the bars at the bottom of the graph are proportional to these deviations. We see that the largest difference between TC1 and the other statistics arises when $|\theta_j - \theta|$ is small; on the other hand, TC1 and TC2 nearly coincide when $|\theta_j - \theta|$ is approximately equal to the between-area variance. The empirical model-based composite estimators also perform poorly in all the areas. These results can be summarized as follows:




Figure 2: Root-MSE (RMSE) of each area when the θ_j are fixed across replications. The within-area distribution is normal, with homocedastic within-area variances. The area sample size is constant and equal to 10. The number of replications is 3000. The legend shows the mean and the median of the RMSE across areas.

- TC1 is the most efficient estimator for all the areas. This is a theoretical estimator, not feasible in applications. It provides a benchmark against which other estimators can be compared or related.
- CML and C are inefficient for all the areas and specially for those with the largest deviations from the centre (large eccentricity).
- For the model-based estimators (C, CML and TC2), using estimated or true values of the parameters makes very little difference. This is not the case for the design-based estimators; just compare TC1 with CA.
- CML performs poorly for all the areas and specially for those that deviate substantively from the centre. The accuracy of CA increases for small or extreme values of eccentricity.

The above difference among estimators can not be appreciated when observing RMSE averaged across areas.

We will see that these results hold in a variety of circumstances, when we vary the sample size, with large or small number of areas, and also with deviation from normality, both in the within-area distributions and in the distribution that generates the fixed realized values of the area means.

3.3 Non-normality and unequal sample sizes n_j

Now we consider the case where the θ_j have been sampled from an asymmetric distribution (and they stay fixed across replications), and the within-area distribution is non-normal (in fact, it is a chi-square distribution with 1 degree of freedom). The sample size ranges from 6 to 45, the number of areas is 40. The number of replications is 3000. The results are shown in Figure 3. For clarity of the graph, since across areas the maximum difference of the RMSE for C and CML is .03, only the RMSE for C is shown.




Figure 3: Root-MSE (RMSE) of each area when the θ_j are fixed across replications. The within-area distribution is non-normal, with homocedastic within-area variances, and area sample size ranging from 6 to 45. Sample size variation is indicated by the thickness of the bars in the x-axis (thicker bar indicating larger sample size). The number of replications is 3000. The legend shows the mean and median of the RMSE across areas.

Figures (2) and (3) show a similar pattern regarding the relative position of the estimators, though peaks are present in Figure (3) due to the variation of sample size across areas (sample size variation is proportional to the thickness of the bars in the x-axis). Note that the peaks correspond to areas with a relatively small sample size. From Figure (3) we conclude

1. The RMSE tends to increase with the eccentricity of the area.
2. TC1 is superior to all the estimators.

3. The feasible estimators C and CML are inefficient for all the areas, especially on those that deviate highly from the overall mean (high eccentricity), and so is the theoretical estimator TC2.
4. CA does not do as badly as C and CML for those areas with low values on eccentricity.
5. As expected, the RMSE tends to decrease with the sample size.

We also computed a version of the empirical composite C that estimates the variance-within σ_ϵ^2 as an (unweighted) mean of the within-area estimates $\hat{\sigma}_{\epsilon_j}^2$ of (16), but the difference in terms of MSE with the standard version of C was negligible.

We found that the true values of MSE computed according to formula (2) are indistinguishable from the (estimated) ones computed with 3000 replications and presented in Figure 3. Figure 4 displays the same graph with true RMSE for the three estimators D, TC1 and TC2. In both figures we see the superiority of the design-based estimators (TC1) over the model-based ones (TC2), not only for some areas that deviate highly from the overall mean, but also for those areas that exhibit a small value of eccentricity (the areas on the left of the x-axis).




Figure 4: Theoretical values of the RMSE for D, TC1 and TC2 and for each area, for θ_j fixed across replications. Area sample size ranging from 6 to 45. The legend shows the mean and the median of the RMSE across areas.

Table 1: Population characteristics

Counties ('Comarques')	Code	Size N_j	Mean θ_j	Squared bias $(\theta_j - \theta)^2$	Variance $\sigma_{j\epsilon}^2$
Alt Camp	AC	1282	8.73	0.09	3250.37
Alt Empordà	AE	4712	5.28	14.11	294.27
Alt Penedès	AP	3052	8.91	0.02	1686.24
Alt Urgell	AU	745	4.71	18.70	158.25
Alta Ribagorça	AR	140	4.59	19.73	205.38
Anoia	AN	3264	7.86	1.37	801.64
Bages	BA	5698	8.24	0.63	1356.90
Baix Camp	BC	5530	6.47	6.59	6479.54
Baix Ebre	BB	2237	6.31	7.41	534.40
Baix Empordà	BE	4634	5.44	12.92	425.17
Baix Llobregat	BL	20541	9.73	0.48	1642.46
Baix Penedès	CP	2197	5.26	14.23	171.82
Barcelonès	BN	88331	10.63	2.55	10314.88
Berguedà	BG	1397	5.44	12.90	196.15
Cerdanya	CR	788	3.71	28.34	71.93
Conca de Barberà	CB	611	8.29	0.56	1388.95
Garraf	GR	3466	6.28	7.62	685.91
Garrigues	GS	516	5.24	14.42	96.89
Garrotxa	GX	1909	7.51	2.33	419.72
Gironès	GI	6369	9.82	0.62	2037.47
Maresme	MA	11718	6.46	6.64	605.07
Montsià	MO	1918	5.61	11.73	246.00
Noguera	NG	1128	5.12	15.30	93.29
Osona	OS	5494	7.09	3.77	774.65
Pallars Jussà	PJ	410	4.37	21.76	130.37
Pallars Sobirà	PS	272	4.06	24.76	55.46
Pla d'Urgell	PU	1106	6.59	5.95	271.85
Pla de l'Estany	PE	1160	6.07	8.79	143.37
Priorat	PR	254	4.11	24.26	180.17
Ribera d'Ebre	RE	620	5.71	11.07	418.72
Ripollès	RI	959	7.87	1.35	875.92
Segarra	SG	594	10.87	3.35	8171.41
Segrià	SR	7096	7.74	1.69	714.23
Selva	SV	4586	7.11	3.70	610.20
Solsonès	SO	508	5.58	11.93	157.58
Tarragonès	TG	7440	9.42	0.15	1675.66
Terra Alta	TA	297	4.25	22.87	40.28
Urgell	UG	1178	6.28	7.59	312.25
Val d'Aran	VA	503	5.28	14.08	270.11
Vallès Occidental	VC	26683	10.34	1.71	3026.89
Vallès Oriental	VR	11795	8.45	0.34	832.68

† The average number of affiliates in Catalonia (overall mean θ) is 9.04.

Figures 3 and 4 show the same ranking among the estimators according to their root-MSE. The same conclusions 1 to 4 that were drawn from Figure 3 apply also to Figure 4.

4 Simulation study on a real population

In this section we study the behaviour of several estimators through a Monte Carlo simulation in which we replicate samples from the Labour Force Census of Enterprises affiliated with the Social Security system in Catalonia. This census contains information on the number of employees who are registered in the Social Security system for each enterprise. The data is available on a quarterly basis from year 1992. We consider only the population in the first quarter of year 2000. The census contains 243,184 observations for Catalonia in year 2000, divided into 12 groups according to the economic sector to which each firm belongs, and into 41 counties (the ‘comarques’).

We ignore the sector-based classification and focus solely on the division by counties. Table 1 shows the number of enterprises (population size) and the mean and variance of the variable of interest (number of registered employees) in each county. The distribution of the enterprises across Catalonia is very uneven, as they are concentrated mainly in densely populated areas. In our set-up, the small areas are held fixed across resampling over the 1000 replications. In each replication, we extract a proportional stratified sample by county. We used sample sizes representing 10%, 5%, 2% and 1% of the population, which gives sample sizes close to those used by IDESCAT in several surveys. Table 2 summarizes the characteristics of these samples. Sample sizes for each county can be easily deduced from Table 1, applying the corresponding sample size percentage reported in each simulation.

Table 2: Sample sizes of the Monte Carlo study (empirical population)

Overall sample		Sample size in county			
% of pop.	Sample size	Mean	Median	Min.	Max.
1	2431	59.3	19	1	883
2	4863	118.6	38	3	1767
5	12159	296.6	95	7	4417
10	24316	593.1	191	14	8833

This population recreates conditions of non-normality, uneven sample sizes, and heterogeneity of within-area variances that are likely to appear in applications. So a Monte Carlo evaluation based on this population will assess the performance of competing estimators in a realistic setting.

We evaluate the performance of the theoretical estimators TC1 and TC2, and the empirical estimators C, CML and CA described in Section 2. For each estimator,

we computed the empirical relative root-MSE (RRMSE) across replications for each county. Using absolute (instead of relative) root-MSE gave the same pattern of performance as when using RRMSE. Even though the theoretical estimators TC1 and TC2 are unfeasible in practice (since the true values of the variances are unknown in a given application), we wanted the Monte Carlo study to illustrate the effect on accuracy when moving from BLUP to EBLUP. The graphs show clearly a shift on accuracy between the theoretical estimators TC1 and TC2 (BLUE estimators, from the fixed and random perspective, respectively) and the other empirical (EBLUP) estimators, with the theoretical estimators having, of course, lower MSE. In this Monte Carlo study, the theoretical estimators should be viewed as providing benchmarks for the accuracy of the EBLUP estimators. While TC1 is the best estimator of them all, its empirical counterpart CA performs worse in average. Also, the empirical C performs worse on average than its theoretical counterpart TC2. The discrepancy between the empirical estimators and their theoretical counterparts is larger when the areas are not too extreme. The CML estimator was computed using `proc xtmixed` of the software package `Stata 10.0`, employing the option `emonly`.⁴ For each estimator, we computed the empirical relative root-MSE (RRMSE) across replications for each county. For the 10% and 5% sample sizes, the direct and the composite estimators have similar RRMSE values. For those sampling schemes, D has the smallest RRMSE among the feasible estimators and is more efficient than the theoretical model-based TC2 estimator. We therefore focus on the description of the 2% and 1% sampling designs.

Figure 5 plots the variation of the RRMSE for the estimators and areas for the 2% sampling design. For clarity, the RRMSE of CML is omitted as it is nearly indistinguishable from C. The same pattern of variation is observed for the 1% sample. Areas have been ordered with respect to their eccentricity, i.e. deviation of the area mean from the overall mean (the heights of the bars are proportional to the eccentricity of the area). We see that the RRMSE tends to increase as the areas become more extreme in terms of eccentricity. The area sample size is proportional to the thickness of the bar. We observe that RRMSE tends to decrease as the sample size increases.

The direct estimator D performs poorly. The design-based estimator CA has a very good performance across all the areas, being close to its theoretical counterpart, TC1, which is the most efficient. The model-based estimators TC2 and C (and CML) do better than the direct estimator D but worse than the theoretical design-based estimator TC1. The poor performance of D and other feasible model-based estimators for some counties, specially for the county SG ('Segarra') stands out. Segarra has both a huge value of the within-county variance and a very small sample size (see columns 3 and 5 of Table 1). In applications of small-area estimation, it should be of high concern that our area has such extreme features. If we knew the true values of the squared area-bias and

⁴ Using `proc xtmixed` with default options produced a large percentage of replications with non-convergent solutions; however, the problem of non-convergence disappeared when we used the option `emonly` (expectation maximization algorithm when the gradient based routines did not converge). That is, a proper CML estimate was obtained in each replication.

the within-area variance (as in TC1 and TC2) then MSE would be reduced dramatically for SG.

The high fluctuation of the performance of the direct estimator is due to the variation of the sample sizes and within-area variances across areas. For extreme areas, the CA estimator performs similarly as the design-based estimator TC1. Estimation of the population parameters has a profound effect on the accuracy for areas that are not extreme. This is the case, for example, for Alt Camp (AC).

For completeness, Figure 6 shows the results for the 10% sampling design. We observe the same pattern of performance across areas as in Figure 5. The distance between the model-based and the design-based estimators is more obvious for areas with greater eccentricity. This graph shows that the direct estimator nearly matches the efficiency of TC1, a clear indication that for such a large sample size, small-area estimation is redundant.

Figures 5 and 6 show that the model-based estimators do not perform too badly when the small areas do not show a high value of eccentricity. But for areas that are very extreme these estimators do worse than the the direct estimator. Figure 5 shows that CA can be a good alternative in practice, because TC1 is unknown in real cases. Surprisingly, not even averaging across areas, the estimator C (or CML) can compete with the composite alternative CA.

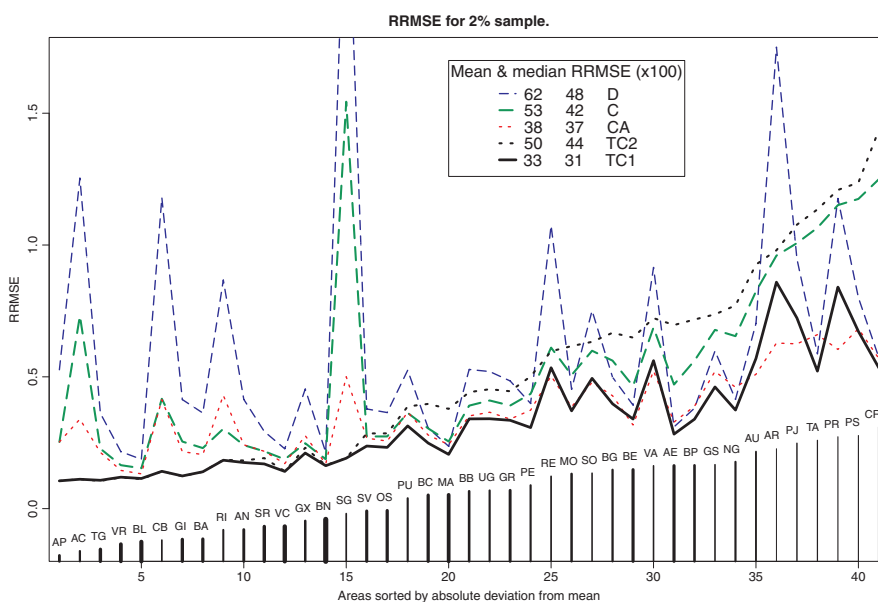


Figure 5: For the Monte Carlo analysis with an empirical population the graph shows the RRMSE or each county, for the 2% sample. Heights of the bars are proportional to the deviations or the area-level means from the overall mean and their thicknesses are proportional to sample sizes. The legend shows the mean and the median of the RRMSE across areas.

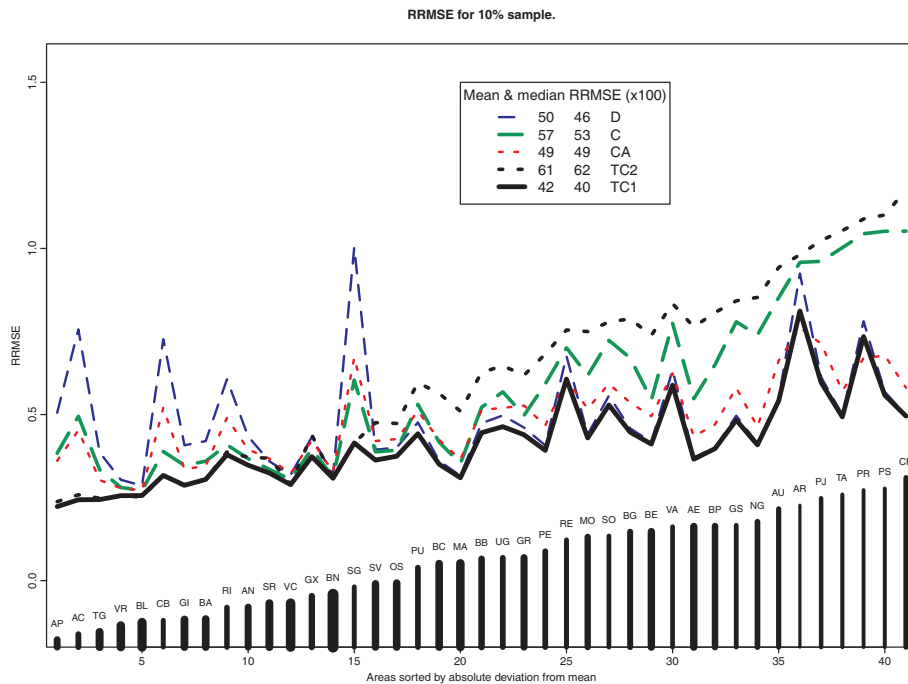


Figure 6: For the Monte Carlo analysis with an empirical population the graph shows the RRMSE for each county, for the 10% sample. Heights of the bars are proportional to the deviations or the area-level means from the overall mean and their thicknesses are proportional to sample sizes. The legend shows the mean and the median of the RRMSE across areas.

This Monte Carlo study with a real population provides a context in which we can recognize different scenarios encountered in an application.

1. For an area with a large sample size, all the small-area estimators are close to each other. This is the case of Barcelonès (BN).
2. In areas with a small sample size and a extreme within-area variance, not necessarily extreme in eccentricity, the empirical small-area estimators may perform very poorly. This is the case of Segarra (SG). For such areas, the incorporation of information on the magnitude of the between- and within-area variances may produce dramatic gains on RRMSE.
3. In an area with a small value of eccentricity and small sample size, model-based estimators are less efficient than the design-based estimators. This can be seen in Alt Camp (AC).
4. In an area with a high value of eccentricity and small sample size, design- and model-based estimators gives high gains over the direct estimator. This can be seen in Alta Ribagorça (AR).

5 Conclusions and agenda

We have seen that in the case of fixed areas, averaging MSE across areas does not provide the complete picture of the performance of alternative small-area estimators. Such averaging of MSEs can be used only to evaluate accuracy in the context of random-area parameters.

We conclude that a composite estimator that uses a common bias estimator for all the areas performs poorly on areas that are extreme. The same is true for the theoretical composite estimator TC2. The problem carries over to the mixed-effects regression, even when the model is not misspecified. Therefore, estimation of the squared bias term for each area becomes crucial.

We conjecture that issues regarding the estimation of the variances within the areas will be less critical; however, the exercise on a real population shows also the importance of recognizing non-normality and variation across areas of the within-area variance.

These findings indicate that the key to improve small-area estimation is to acknowledge the fixed-effect nature of the data and to improve estimation of the squared area-bias. Differences (heteroscedasticity) of the within-area variances seem to be also critical. Several alternatives arise:

1. Using auxiliary information (such as a census or a previous survey) to estimate the squared bias. Then the same simple and convenient composite estimators could be used.
2. Improving the alternative composite estimator by defining different groups of areas that share a common between-area variance.
3. Estimating the squared bias using small-area methods. This approach has already been used in Longford (2007) for estimating MSEs of model-based estimators.

Further work assessing these alternatives is needed. As a final remark, we should note that we have confined discussion to the most simple model set-up where covariates are not present in the model; additional work assessing the validity of our findings when the model is expanded is worth pursuing.

References

- Costa, A., Satorra, A. and Ventura, E. (2003). An empirical evaluation of small area estimators. *SORT (Statistics and Operations Research Transactions)*, 27, 113-135.
- Efron, B., and Morris C. E. (1973). Stein's estimation rule and its competitors – an empirical Bayes approach, *Journal of the American Statistical Association*, 68, 117-130.
- Ghosh M. and Rao, J. N. K. (1994). Small area estimation: an appraisal, *Statistical Science*, 9, 55-93.

- Isaki, C. T. (1990). Small-area estimation of economic statistics, *Journal of Business & Economic Statistics*, 8, 435-41.
- Longford, N. T. (2007). On standard errors of model-based small-area estimators, *Survey Methodology*, 33, 69-79.
- Neudecker, H. and Satorra, A. (2003). On best affine prediction, *Statistical Papers*, 44, 257-266.
- Platek, R., Rao, J. N. K., Särndal, C. E. and Singh, M. P. (Eds.) (1987). *Small Area Statistics: An International Symposium*, John Wiley and Sons: New York
- Rao, J. N. K. (2003). *Small Area Estimation*. John Wiley and Sons: New York
- Santamaría, L., Morales, D. and Molina, I. (2004). A comparative study of small area estimators, *SORT*, 28, 215-230.
- Singh, M. P., Gambino, J. and Mantel H. J. (1994). Issues and strategies for small area data, *Survey Methodology*, 20, Statistics Canada, 3-22.