Small Area Estimates of Crime Rates for States and Large Counties Based on the NCVS

Robert E. Fay[1] and Mamadou S. Diallo[1]
[1]Westat, Inc., 1600 Research Boulevard, Rockville, MD 20850

Abstract
The National Crime Victimization Survey (NCVS) in the U.S. has provided estimates of violent and property crime for over four decades. Until recently, the survey has had almost exclusively a national focus. As recommended by a National Academy of Sciences panel, the Bureau of Justice Statistics (BJS) has undertaken a number of efforts to expand the geographic utility of the survey results. This paper is an outgrowth of research to provide small area estimates of key crime rates from the NCVS for states, large counties, and large metropolitan areas. The estimates are based a modified version of a time-series model proposed by Rao and Yu to take advantage of strong area-level correlations in the estimates over time. A multivariate version of the model was used to provide estimates for components of the crime rates by type of crime and by relationship to the perpetrator. BJS plans to make the estimates available to users on their website. This paper describes a hybrid model representing a further extension. These methods have potential applications to other situations in which the underlying characteristic exhibits strong stability over time.

Key Words: EBLUP, Rao-Yu model, Fay-Herriot model.

1. Introduction

The Bureau of Justice Statistics (BJS) sponsors and analyzes the National Crime Victimization Survey (NCVS), an ongoing national survey of the civilian, non-institutional population in the U.S, ages 12 and over. The survey produces annual estimates of crime as reported by its victims. Although the survey collects substantial detail about the perpetrator or perpetrators, location, and role of the police (if any), a small number of core variables are of central interest. BJS publishes overall rates for violent crime and property crime without combining them into a single measure. Violent crimes are broadly classified into the categories of rape and sexual assault, aggravated assault, simple assault, and robbery. With the exception of the proxy responses allowed under restricted rules, violent crimes are measured through self-response, with NCVS interviewers contacting individual household members twice a year to interview them about violent crime incidents. A single household respondent is asked about property crimes, which are classified as burglary, motor vehicle theft, and other theft. The Census Bureau conducts the survey for BJS.

In the past, the geographic detail available from the survey has been limited. But in recent years, BJS has investigated alternative approaches to provide more information on the geographic distribution of crime, including funding a recent boost of the NCVS sample in a select set of large states. As part of the effort to provide geographic detail, BJS funded a grant to Westat supporting research to develop small area estimation approaches for the

NCVS. By the end of our grant in December, 2014, we provided BJS sets of small area estimates for states, large counties, and large metropolitan areas.

The small area estimates are in the form of 13 overlapping sets of 3-year average crime rates over a 15-year period for each characteristic. Violent crime estimates are available by type of crime: (1) robbery; (2) aggravated assault, sexual assault, and rape; and (3) simple assault. Violent crime is also available by relationship to the perpetrator: (1) stranger or strangers; (2) intimate partner; and (3) others including other family members, friends and acquaintances. Property crime is disaggregated into (1) burglary, (2) motor vehicle theft, and (3) all other theft. Estimates for states cover 1999-2013. We also provided estimates in the same form for the 65 largest counties and for the 51 largest metropolitan areas for the period 1998-2012. At the time of this writing, BJS is determining the format and manner of their public release.

Our approach to developing the small area estimates was shaped by major features of the application. The NCVS sample design is a multi-stage sample of housing units and non-institutional group quarters. For over two decades, the survey has produced estimates of victimization for the same major categories of crime by largely consistent methods. Evidence from FBI crime statistics also indicates that the relative geographic distribution of crime is highly stable over time. These circumstances suggested that a small area approach combining information across time would outperform the use of cross-sectional small area models.

Basic aspects of the NCVS sample design constrain the choice of models. In particular, non-zero correlations between the sample estimates are present. While some small area models assume statistically independent sampling errors across time, the NCVS design results in sampling correlations for two reasons. First, the survey employs a rotating panel design, with sampled households remaining in sample for a total of 7 interviews scheduled 6 months apart. Respondents report on crimes during the previous 6 months. Households who move from a sampled housing unit are not followed, but instead they are replaced by incoming households. But a substantial portion of the sample is interviewed repeatedly, resulting in a sampling correlation to the extent that some sample individuals are at higher risk of being victims of crime than others. Second, the first-stage design of the survey historically has been redesigned on a ten-year cycle, so most adjacent years are estimated based on the same first-stage sample of primary sampling units, again resulting in correlation across time.

Of the models reviewed by Rao (2003) appropriate to an annual time series of observations, the models proposed by Rao and Yu (1992, 1994) appeared the most applicable to the NCVS situation (Li, Diallo, and Fay, 2012). But the original formulation of the Rao-Yu model assumed stationarity, which appeared contradicted by FBI crime data. Fay and Diallo (2012) introduced a modification to the Rao-Yu model, which we termed the *dynamic model*. Analysis of FBI data showed a marginally better fit with the dynamic model than the Rao-Yu model. Consequently, we used the dynamic model to produce the small area estimates for states, large counties, and large metropolitan areas.

A primary purpose of this paper is to examine a new extension of the Rao-Yu model, which combines the features of the dynamic and Rao-Yu models into a *hybrid model.* In the univariate case, the hybrid model has two more parameters than either the Rao-Yu or dynamic model and can produce the fits of either. Another extension studied by Diallo

(2014) in effect removes the condition of stationarity from the Rao-Yu model, but this extension is not covered here.

We propose the hybrid model as a possible improvement on the dynamic model, but we are unable to report on its application to the NCVS data at this point. Instead, this paper introduces the model and compares it to the Rao-Yu and dynamic models. The hybrid model clarifies the relationship between the two other models and enables a data-driven choice between them.

A previous paper jointly authored with Michael Planty of BJS (Fay, Planty, and Diallo, 2013) reviewed earlier work (Li, Diallo and Fay, 2012; Fay and Diallo, 2012) to develop small area models for the NCVS. The three papers provide more background on the NCVS itself and summarize the specific features of the NCVS small area application leading to the selection of the model and predictor variables. The models use data from the FBI's Uniform Crime Reports (UCR) as auxiliary information, and the previous papers describe this source of data more fully. We refer to the previous papers for specifics of the NCVS application.

Section 2 will review the Rao-Yu, dynamic, and hybrid models. Section 3 presents simulation results analyzing the performance of the models in situations where the underlying model may not be correct. Section 4 comments on the implementation of the hybrid model in R and discusses the potential usefulness of the new model and future research questions.

2. The Rao-Yu Model and Its Generalizations

2.1 The Rao-Yu Model

The Rao-Yu model (1992, 1994) was summarized by Rao (2003). It builds on a linear mixed model for the population values, θ_{it}, t=1,…,T,

$$\theta_{it} = x_{it}'\beta + v_{(ry)i} + u_{(ry)it}$$

where

x_{it}' is a row vector of known auxiliary variables,
β is a vector of fixed effects
$v_{(ry)i}$ is a random effect for area i, $v_{(ry)i} \sim^{iid} N(0,\ \sigma^2_{(ry)v})$
$u_{(ry)it}$ is a random effect for area i, time t, with
$u_{(ry)it} = \rho_{(ry)} u_{(ry)i,t-1} + \epsilon_{(ry)it}$, $|\rho_{(ry)}| < 1$ and $\epsilon_{(ry)it} \sim^{iid} N(0, \sigma^2_{(ry)})$,

In the original version of the model as shown above, the $u_{(ry)it}$´s are assumed to form a stationary time series ($|\rho_{(ry)}| < 1$). The model for the observed sample values, y_{it}, is

$$y_{it} = \theta_{it} + e_{it} = x_{it}'\beta + v_{(ry)i} + u_{(ry)it} + e_{it}$$

where

e_{it} is random sampling error for area i, time t, with
$\mathbf{e_i} = (e_{i1}, \ldots, e_{iT})' \sim N_T(0, \mathbf{\Sigma}_i)$, and where $\nu_{(ry)i}$, $\epsilon_{(ry)it,}$ and $\mathbf{e_i}$ are mutually

independent.

As previously noted, $\mathbf{\Sigma}_i$ need not be diagonal, that is, the model can accommodate sampling covariance across time.

Rao and Yu (1992, 1994) derived the best linear unbiased predictor (BLUP) for this model for known values of the variance parameters $\sigma^2_{(ry)v}$, $\sigma^2_{(ry)}$, and $\rho_{(ry)}$. Adapting a method of moments approach developed by economists to estimate these parameters, they proposed a resulting empirical best linear unbiased predictor (EBLUP). But they noted considerable difficulty in applying the method because of the instability in estimating $\rho_{(ry)}$. They also investigated an EBLUP based on estimating $\sigma^2_{(ry)v}$ and $\sigma^2_{(ry)}$ given a presumed value of $\rho_{(ry)}$. Rao (2003) summarized these results, again noting practical difficulties with the estimation of the variance parameters.

3.2 The Dynamic Model

Fay and Diallo (2012) noted that although the Rao-Yu model might provide a reasonable summary of the UCR crime data, the stationarity assumption appeared somewhat questionable. We proposed instead a minor modification to the Rao-Yu model that would remove the stationarity requirement by modifying the random effect terms. The mixed model for the population values is

$$\theta_{it} = x'_{it}\beta + \rho^{t-1}_{(dyn)} v_{(dyn)i} + u_{(dyn)it};$$

where

$v_{(dyn)i} \sim^{iid} N(0, \sigma^2_{(dyn)v})$ is a random area effect for area i at time $t = 1$,
$u_{(dyn)i1} = 0$, and
$u_{(dyn)it} = \rho_{(dyn)} u_{(dyn)i,t-1} + \epsilon_{(dyn)it}$, for $t > 1$, where
$\epsilon_{(dyn)it} \sim^{iid} N(0, \sigma^2_{(dyn)})$.

The sampling model is modified similarly

$$y_{it} = \theta_{it} + e_{it} = x'_{it}\beta + \rho^{t-1}_{(dyn)} v_{(dyn)i} + u_{(dyn)it} + e_{it}$$

where again

e_{it} is random sampling error for area i, time t, with
$\mathbf{e_i} = (e_{i1}, \ldots, e_{iT})' \sim N_T(0, \mathbf{\Sigma}_i)$, and where $\nu_{(dyn)i}$, $\epsilon_{(dyn)it,}$ and $\mathbf{e_i}$ are mutually independent.

Unlike $\rho_{(ry)}$ in the original Rao-Yu model, $\rho_{(dyn)}$ is not constrained to be less than 1. When $\rho_{(dyn)} > 1$, the model corresponds to a situation in which areas progressively diverge. When $\sigma^2_{(dyn)v} = \sigma^2_{(dyn)}/(1 - \rho^2_{(dyn)})$ and $|\rho_{(dyn)}| < 1$, the dynamic model becomes equivalent a Rao-Yu model with $\sigma^2_{(ry)v} = 0$, $\sigma^2_{(ry)} = \sigma^2_{(dyn)}$ and $\rho_{(ry)} = \rho_{(dyn)}$. But by dropping the stationarity assumption, the dynamic model is more appropriate for a situation in which the disparity among states dissipates over time.

Fay and Diallo (2012) compared the fit of the Rao-Yu model and the dynamic model to the UCR data at the state level, using a model that included only main effects for each year. For all of UCR variables, the dynamic model provided a better fit than the Rao-Yu model, in most cases by a statistically significant amount. Nonetheless, considering the size of the UCR data set, the improvements in fit were modest, suggesting that the Rao-Yu model would have provided an adequate alternative for the small area estimation application to NCVS.

Fay and Diallo (2012) also reported successfully developing maximum likelihood (ML) and restricted maximum likelihood (REML) approaches to estimating the variance parameters, including ρ, both for the dynamic model and for the original Rao-Yu model, simply by following the general approach summarized by Rao (2003, Section 6.2) for the general linear mixed model. The possible use of ML and REML in this context was likely subsequently discovered by other researchers; for example, the REML approach was used for a spatial-temporal model that generalizes the Rao-Yu model (Marhuenda, Molina, and Morales, 2013).

3.3 The Hybrid Model

The hybrid model combines features of the dynamic and Rao-Yu models. The form of the model is:

$$\theta_{\mathrm{it}} = \mathrm{x}'_{\mathrm{it}}\beta + \rho_{(dyn)}^{t-1} v_{(dyn)i} + u_{(dyn)it} + u_{(ry)it};$$

where, as before

$v_{(dyn)i} \sim^{\mathrm{iid}} \mathrm{N}\left(0, \sigma^2_{(\mathrm{dyn})\mathrm{v}}\right)$ is a random area effect for area i at time $t = 1$,
$u_{(dyn)i1} = 0$,
$u_{(dyn)it} = \rho_{(dyn)} u_{(dyn)i,t-1} + \epsilon_{(\mathrm{dyn})\mathrm{it}}$, for $\mathrm{t} > 1$, where
$\epsilon_{(\mathrm{dyn})\mathrm{it}} \sim^{iid} N\left(0, \sigma^2_{(dyn)}\right)$
$\mathrm{u}_{(\mathrm{ry})\mathrm{it}}$ is a second random effect for area i, time t, with
$\mathrm{u}_{(\mathrm{ry})\mathrm{it}} = \rho_{(ry)} \mathrm{u}_{(\mathrm{ry})\mathrm{i,t-1}} + \epsilon_{(\mathrm{ry})\mathrm{it}}$, $\left|\rho_{(ry)}\right| < 1$ and $\epsilon_{(\mathrm{ry})\mathrm{it}} \sim^{\mathrm{iid}} N\left(0, \sigma^2_{(ry)}\right)$.

For estimability, the model additionally requires $\rho_{(ry)} < \rho_{(dyn)}$. In practice, a minimum difference of the two correlation parameters can be set and enforced during iteration. When the estimated $\sigma^2_{(ry)}$ = 0, $\rho_{(ry)}$ is not estimable, and the hybrid model reduces to the dynamic model.

Finding the maximum of the restricted likelihood for the hybrid model is challenging because the restricted likelihood surface often has local maxima. The current version of the algorithm implemented here considers two starting values, one based on the REML parameter estimates for the Rao-Yu model and the other for the dynamic model. The parameters for the two starting values are iterated 15 times under the restricted likelihood for the hybrid model. The parameter set with the higher restricted likelihood is then iterated to convergence. The algorithm is not guaranteed to find the global maximum of the restricted likelihood, but as a consequence of its starting values it provides a fit with a greater restricted likelihood than either the Rao-Yu or dynamic fits of the data.

3.4 Multivariate Version of the Hybrid Model

Fay, Planty, and Diallo (2013) presented a multivariate version of the dynamic model. The multivariate version was motivated by the objective of jointly modeling the components of crime and their sum. For example, the goal was to obtain small area estimates of burglary, motor vehicle theft, and other theft as the components of total property crime, which was also to be estimated. The multivariate approach addresses the problem that univariate modeling of each component and their sum separately would produce a set of inconsistent estimates. The BLUP estimator is, by definition, a linear function of the observed **y,** and the results for BLUP for the general linear mixed model provides simultaneously the BLUP for any linear combination of the fixed and random effects (Rao, 2003, Section 6.2.1). As a consequence, the BLUP for the sum of crime rates by type of crime is the sum of the BLUPs for the components. (In our application, the result applies equally to 3-year averages of crime rates.) Rao (2003, section 8.1) reviewed a number of earlier applications of multivariate models to small area estimation, although they remain less frequently used than univariate models.

Similarly, a multivariate version of the hybrid model may be developed. The population values for area i, time t, can be represented as a vector $\boldsymbol{\theta}_{it} = (\theta_{it1}, \theta_{it2,} \ldots, \theta_{itp})'$. By letting k index the components of the multivariate vector, a model for the population values can be expressed

$$\theta_{\mathrm{itk}} = \mathrm{x}'_{\mathrm{itk}}\beta_k + \rho_{(dyn)}^{t-1} v_{(dyn)ik} + u_{(dyn)itk} + u_{(ry)itk};$$

for $k = 1, \ldots, p$, where

$\boldsymbol{v}_{(\mathrm{dyn})\mathrm{i}} = (v_{(dyn)i1}, v_{(dyn)i2}, \ldots)' \sim^{iid} N\big(\mathbf{0}, \boldsymbol{\Sigma}_{(dyn)v}\big)$ is a vector of random effects for area *i* at time *t*=1,

$u_{(dyn)i1k} = 0$ at time *t*=1,

$u_{(dyn)itk} = \rho_{(dyn)} u_{(dyn)i,(t-1),k} + \epsilon_{(\mathrm{dyn})\mathrm{it}k}$, for t > 1, where

$\boldsymbol{\epsilon}_{(\mathrm{dyn})\mathrm{it}} = (\epsilon_{(\mathrm{dyn})\mathrm{it1}}, \epsilon_{(\mathrm{dyn})\mathrm{it2}} \ldots)' \sim^{\mathrm{iid}} N\big(\mathbf{0}, \boldsymbol{\Sigma}_{(\boldsymbol{dyn})}\big)$

$u_{(ry)itk} = \rho_{(ry)} u_{(ry)i,(t-1),k} + \epsilon_{(\mathrm{ry})\mathrm{it}k}$, for t > 1, where

$\boldsymbol{\epsilon}_{(\mathrm{ry})\mathrm{it}} = (\epsilon_{(\mathrm{ry})\mathrm{it1}}, \epsilon_{(\mathrm{ry})\mathrm{it2}} \ldots)' \sim^{\mathrm{iid}} N\big(\mathbf{0}, \boldsymbol{\Sigma}_{(\boldsymbol{ry})}\big)$

and where the series $\mathbf{u}_{(ry)it} = \big(u_{(ry)it1,} u_{(ry)it2}, \ldots\big)$ is stationary. The series $\mathbf{u}_{(dyn)it}$, $\boldsymbol{\Sigma}_{(dyn)v}$ and $\boldsymbol{\Sigma}_{(dyn)}$ are assumed to be related to each other by

$$\Sigma_{(dyn)v(k,k')} = \sigma^2_{(dyn)vk} \mathrm{R}_{(dyn)(k,k')}\, \sigma^2_{(dyn)vk'} \quad \text{and}$$
$$\Sigma_{(dyn)(k,k')} = \sigma^2_{(dyn)k} \mathrm{R}_{(dyn)(k,k')}\, \sigma^2_{(dyn)k'},$$

for a common correlation matrix, $\mathbf{R}_{(\boldsymbol{dyn})}$**;** where $\boldsymbol{\sigma}^2_{(dyn)v} = (\sigma^2_{(dyn)v1}, \sigma^2_{v(dyn)2}, \ldots)'$; and $\boldsymbol{\sigma}^{\mathbf{2}}_{(dyn)} = (\sigma^2_{(dyn)1}, \sigma^2_{(dyn)2} \ldots)'$. Similarly,

$$\Sigma_{(ry)(k,k')} = \sigma^2_{(ry)k} \mathrm{R}_{(ry)(k,k')}\, \sigma^2_{(ry)k'},$$

Note that the proposed model posits only two correlation parameters, $\rho_{(dyn)}$ and $\rho_{(ry)}$, to express the correlation over time. For example, $\rho_{(dyn)} = .97$ and $\rho_{(ry)} = .40$ would mix a dynamic model with high correlation over time with a stationary series with lower correlation over time.

The sampling model is

$$y_{itk} = \theta_{\text{itk}} + e_{itk} = \text{x}'_{\text{itk}}\beta_k + \rho_{(dyn)}^{t-1} v_{(dyn)ik} + u_{(dyn)itk} + u_{(ry)itk} + e_{itk}$$

for $k = 1, \dots, p$, where $\mathbf{v}_{(dyn)i}$, $\boldsymbol{\epsilon}_{(dyn)it,}$ and $\boldsymbol{e_i} = (e_{i11,} e_{i21}, \dots, e_{iT1}, e_{i21,} \dots)'$ are mutually independent.

Again, existing theory (Rao, 2003, section 6.2) provides the general mathematical results to implement this model, including ML and REML estimation of the parameters and mean square error estimation for the REML results based on extensions of the methods begun by Prasad and Rao (1990).

3. Simulation Comparisons

3.1 Setup

A simulation was designed to be similar to the NCVS application, where the sample sizes in the target areas varied substantially. The simulation used 48 areas divided into 4 groups of 12 areas each. The simulation results were assessed by averaging performance within each of the 4 groups. For simplicity, the sampling covariance matrix over time was assumed to be a multiple of the identity matrix for each area. Two sample sizes were considered: one resulting in sampling variances for each year of .1, .2, .4, and .8 for each of the groups, respectively; and a second with sampling variances for each year of .3, .6, 1.2, and 2.4. These variance choices can be thought of as “large” and “small” sample sizes for the overall survey, respectively, in that the first set of variances could correspond to the sampling errors of estimated survey means from a large sample and the second a set to a small sample. A series of 15 years was generated in each case. Each simulation was replicated 1,000 times.

Five hybrid parameter conditions were studied. For the first three, only the hybrid model fully fit the expected values of the data (Table 1). The fourth condition was consistent with both the hybrid model and the dynamic model but not the Rao-Yu model. Conversely, the fifth condition was consistent with the hybrid model and the Rao-Yu model but not the dynamic model.

Table 1: Parameter Choices for Five Simulation Conditions

	Condition 1	*Condition 2*	*Condition 3*	*Condition 4*	*Condition 5*
$\sigma^2_{(\text{dyn})\text{v}}$	1.00	1.00	1.00	1.16	1.00
$\sigma^2_{(\text{dyn})}$	.01	0	0	.02	0
$\rho_{(dyn)}$	.97	.97	.90	.97	1.00
$\sigma^2_{(\text{ry})}$	.15	.15	.15	0	.13
$\rho_{(ry)}$	.40	.40	.40	0	.80

The true population means were generated according to the random effect parameters, and then sampling errors were added. The small area estimates were compared to the generated true population means. For each pairing of condition and sample size, 1,000 samples were run with normal random errors, and the same data were fitted with the three small area models and compared. The parameters were estimated with REML. Although other properties of the small area models can be studied, the analysis here focuses on the predictive accuracy in the final year of the 15-year series, because users are typically more interested in the most recent value rather than the historical series.

3.2 MSE Results

Under each of the simulation conditions, the small area estimates substantially improve the direct estimates. The mean square errors (MSEs) of the small area estimates represented substantial gains over the direct sample estimates (Fig. 1). The improvements range from about 30% to more than 90%, with greater relative improvements realized for the smaller area groups and small sample.

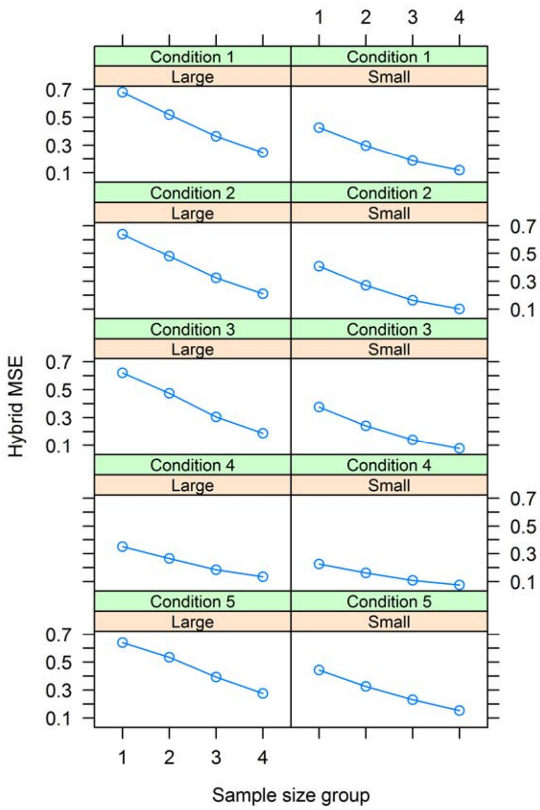

Figure 1: Ratios of the MSE of the small area estimates with the hybrid model to the sampling variance of the direct estimate for the four sample size groups with the simulated large and small samples.

When the MSEs of the three models are compared for the first three conditions, the overall result is that choosing the correct hybrid model over the other two models results in MSE improvements of generally 5% to 20% (Figure 2). Although the Rao-Yu model approaches the performance of the hybrid model under conditions 1 and 2, condition 3 challenges the Rao-Yu model. The dynamic model is a correct model under condition 4, and here a slight advantage can be detected over the hybrid model. Similarly, the Rao-Yu model is best under condition 5, again by a small amount compared to the hybrid model.

The findings confirm general expectations. In large enough samples, it is better to work with the correct model than an incorrect model, but when two models are correct, with one nested inside the other, using the more parsimonious one reduces the variance and consequently the mean square error. Overall, the gains of the hybrid model are enough to consider it as a candidate in situations where either of the two models are applied.

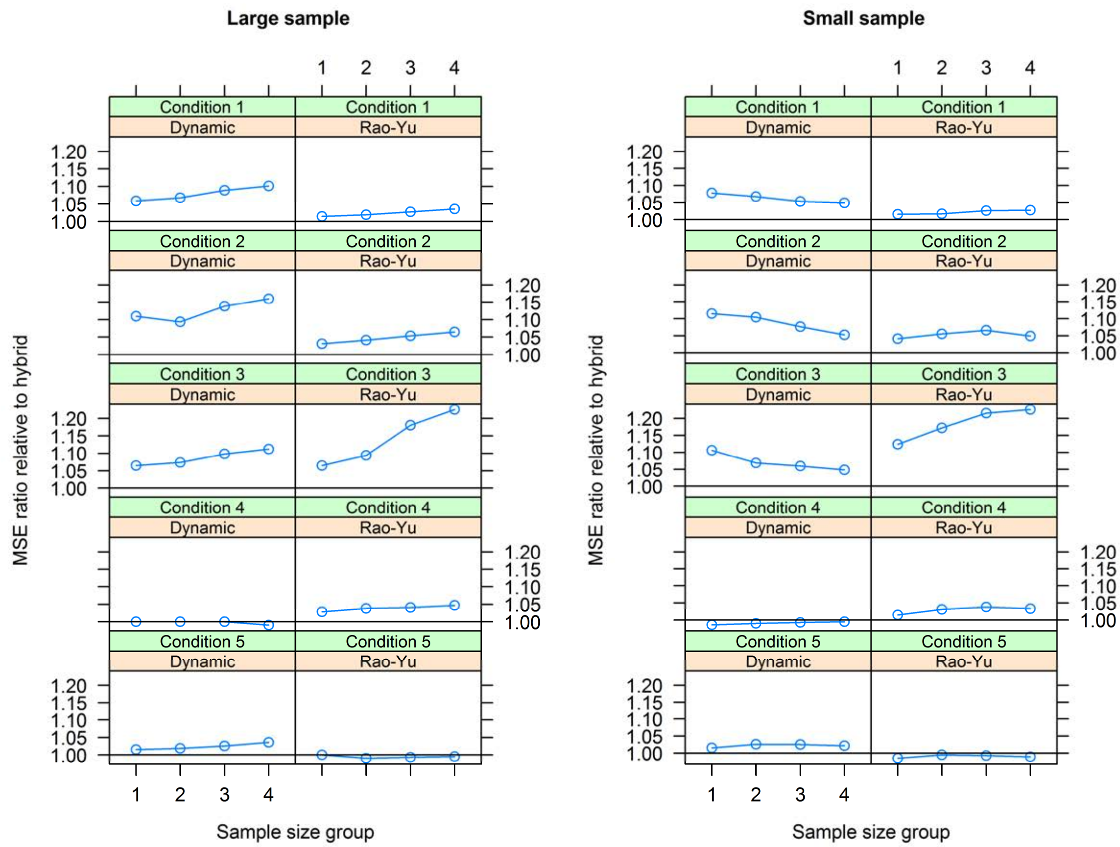

Figure 2: Ratios of the MSE of the small area estimates for the dynamic and Rao-Yu models relative to the hybrid model for the simulated large and small samples.

3.3 Performance of the MSE Estimators

We compared the performance of analytic MSE estimators from existing theory (Rao, 2003, section 6.2) to the MSEs obtained by simulation. When we previously studied the bias of the MSE estimator under the dynamic model, we had found very low bias (e.g., Fay, Planty, and Diallo, 2013). We again found this to be the case for many but not all of the situations studied in the simulation (Figures 3 and 4). The most striking exception appeared for the hybrid model under condition 4, the condition for which both the hybrid and dynamic model fit the data. In this case, the upward bias of the MSE estimator for the hybrid model is about 60% for the large sample and over 250% for the small sample. For the same condition, the MSE estimator for the dynamic model is nearly unbiased, consistent with earlier findings, and the MSE estimator for the Rao-Yu model exhibits a modest upward bias.

Other features of the results are worth noting. For all of the conditions studied, the MSE estimator for the Rao-Yu model is nearly unbiased or has a small upward bias. For the first three conditions, when the dynamic model is not a true model for the population, the MSE estimator shows a downward bias in most cases. The MSE estimator tends to upward bias, particularly for condition 5 as well as condition 4, and for the small sample size.

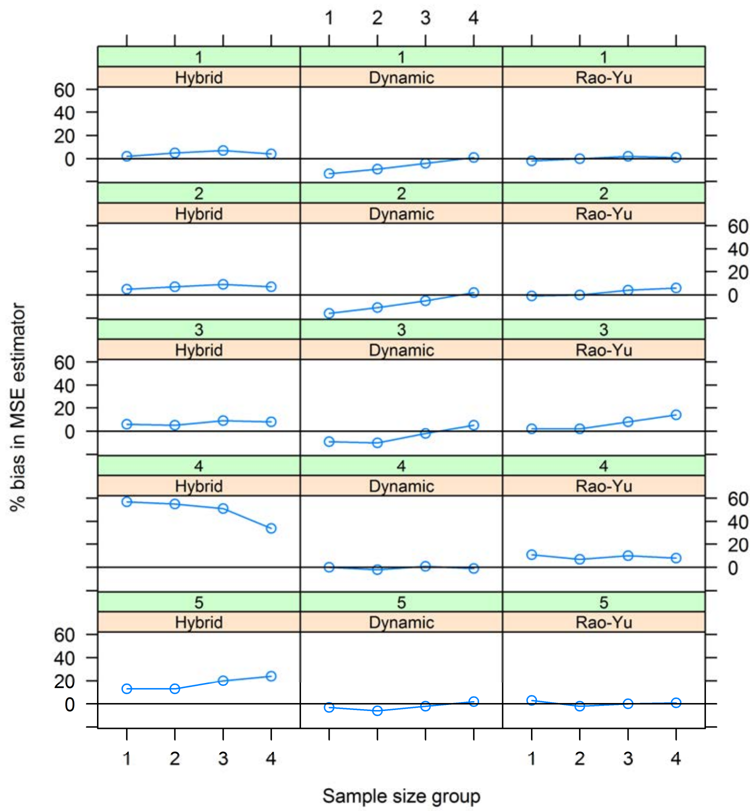

Figure 3: Percent bias in the estimated MSE of the small area estimates for the hybrid, dynamic, and Rao-Yu models compared to the actual MSE estimated from the simulation for the large sample.).

4. Conclusions

In January 2015, the sae2 package was accepted as part of the publicly available CRAN library of R (R Core Team, 2015). A modified version of the package was used for the simulation study, including a new function to compute REML estimates for the hybrid model. After some additional refinements, including support for maximum likelihood estimation, our current goal is to submit a revised version of the package to CRAN before the end of 2015.

The simulation result point to the possibility of further refinements and additions. The analytic approximation to the MSE had appeared satisfactory until the results for condition 4. An alternative approach is to implement a bootstrap estimation of the mean square error, which may hold promise in this case. Other researchers have successfully applied the bootstrap to other small area estimation problems, and some of the packages in R, including sae, implement the bootstrap for MSE estimation for the models they support.

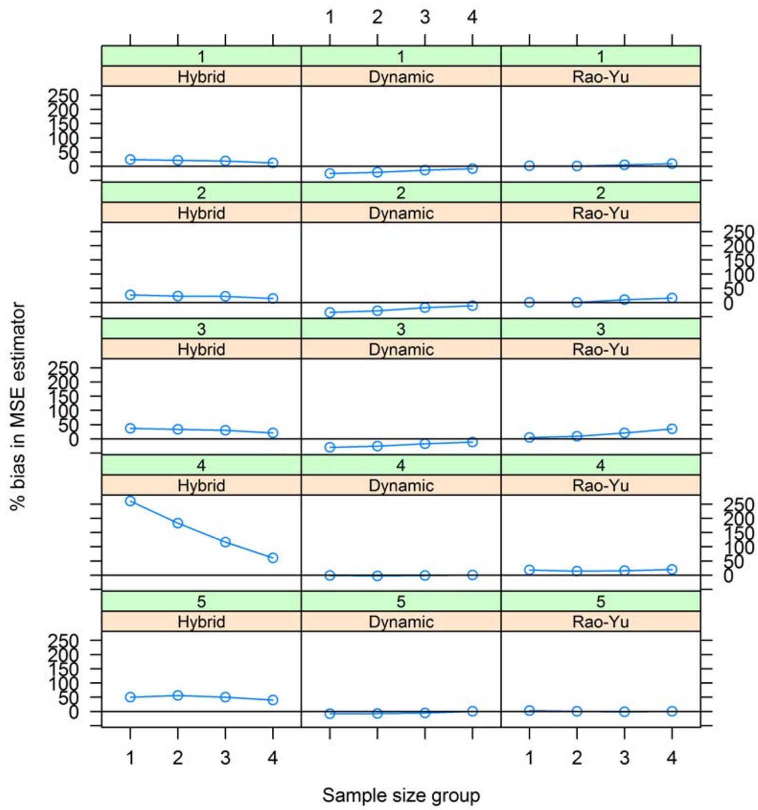

Acknowledgements

The results reflect the views of the authors and not necessarily those of Westat, the Bureau of Justice Statistics, or the Census Bureau. We thank Graham Kalton for his comments on the paper.

References

Diallo, M.S. (2014), "Small Area Estimation under Skew-Normal Nested Errors Model", Ph.D. Dissertation, School of Mathematics and Statistics, Carleton University, Ottawa, Canada.

Fay, R.E. and Diallo, M.S. (2012), "Small Area Estimation Alternatives for the National Crime Victimization Survey," *Proceedings of the Survey Research Methods Section,* American Statistical Association, Alexandria, VA, pp. 3742-3756.

Fay, R.E., Planty, M., and Diallo, M.S. (2013), "Small Area Estimates from the National Crime Victimization Survey," *Proceedings of the Survey Research Methods Section,* American Statistical Association, Alexandria, VA, pp. 1544-1557.

Li, J., Diallo, M.S., and Fay, R.E., (2012) "Rethinking the NCVS: Small Area Approaches to Estimating Crime," presented at the 2012 Federal Committee on Statistical Methodology Conference, Washington, DC, Jan. 10-12, 2012, available at http://www.fcsm.gov/12papers/Li_2012FCSM_I-B.pdf.

Marhuenda, Y., Molina, I., and Morales, D. (2013), "Small Area Estimation with spatio-temporal Fay-Herriot Models," *Computational Statistics and Data Analysis*, 58, 308-325.

Prasad, N.G.N., and Rao, J.N.K. (1990), "The Estimation of Mean Squared Error of Small-Area Estimators," *Journal of the American Statistical Association*, 85, 163-171.

R Core Team (2015), "R: A language and environment for statistical computing," R Foundation for Statistical Computing, Vienna, Austria, URL http://www.R-project.org/.

Rao, J.N.K. (2003), *Small Area Estimation,* John Wiley & Sons, Hoboken, NJ.

Rao, J.N.K. and Yu, M. (1992), "Small Area Estimation Combining Time Series and Cross-Sectional Data," *Proceedings of the Survey Research Methods Section,* American Statistical Association, Alexandria, VA, pp. 1-9.

______ (1994), "Small Area Estimation by Combining Time Series and Cross-Sectional Data," *Canadian Journal of Statistics,* 22, 511-528.