

Assessing Fit of SAIPE Models to Census and CPS County Child-Poverty Rates

Eric V. Slud

Census Bureau, SRD, and Univ. of Maryland
Washington DC 20233-0001 , Eric.V.Slud@census.gov

Abstract. The Small Area Income and Poverty Estimation (SAIPE) project is an ongoing Census effort to estimate numbers of poor school-age children by state, county, and ultimately school district, based upon data from the Current Population Survey (CPS), IRS, Food Stamps, and the latest decennial census. The estimation methodology used at the county level, based on a Fay-Herriot (1979) model fitted to log-counts of related school-age children in CPS-sampled households (for an aggregate of 3 years), discards data from those sampled counties with no sampled poor children. It has been the subject of extensive development at the Census and evaluation by a NAS panel (NRC Report, Citro and Kalton eds., 2000) with respect to ‘internal’ criteria of fit to weighted CPS child-poverty estimates and ‘external’ criteria of fit to decennial Census child-poverty estimates. Further data-analytic evaluations by Maiti and Slud (2002) compared linear model fits (to the log child-poverty rates) with several generalized linear model (GLMM, i.e., logistic regression with random intercept) fits, to the 1990 and 1994 SAIPE county data, and found that the models which fitted best to the CPS estimates did not do so with respect to the Census, and that the quality of fit differed between counties with and without sampled poor children. This paper further studies differences and mutual predictability between Census and CPS county child-poverty rates using additional years of SAIPE data.

Key words: Fay-Herriot model, weighted linear regression, small area estimation.

Acknowledgment. This work relies on the multi-year SAIPE data files assembled in a uniform format by Jerry Maples. This paper describes research and analysis of its author, and is released to inform interested parties and encourage discussion. Results and conclusions are the author’s and have not been endorsed by the Census Bureau.

1. Introduction: SAIPE Goals & Models

Under the terms of Title I of the No Child Left Behind Act, more than \$14 billion in compensatory education funds annually are allocated to counties and school districts using a formula involving child poverty-rate estimates. The SAIPE approach to county-level estimates was developed in response to legislation in 1994 (NRC Report of National Academy of Sciences Panel of Estimates for Small Geographic Areas 2000, p. 3) calling for the Census Bureau to supply ‘updated estimates’ of county-level child poverty for use in Title I allocations to counties in 1997-98 and 1998-99, and thereafter to provide estimates at school-district level. Estimates were to be based on a SAIPE county *small area estimation* model using decennial-census and administrative predictor variables to express the similarity of CPS child-poverty data across counties, thereby ‘borrowing strength’ (Ghosh and Rao 1994) from observed data to compensate for the absence of many counties from 3-year CPS samples and for the smallness of samples in many other counties. (CPS sample data were aggregated for stability over three years, including the year before and the year after each income-year of interest). The CPS, the primary national survey measuring population and poverty each year, provides the SAIPE program with national county-level child poverty estimates, through sample-weighted estimates of numbers and proportion of poor children among children aged 5–17 related to primary householder (*poor related school-age children*). The administrative predictors are: the county numbers of IRS child exemptions for families in poverty and of all child exemptions reported on tax returns, along with county numbers of households (HUs) participating in the Food Stamp program.

The modeling framework in the SAIPE program (NRC Report 2000, Appendix A by W. Bell) is that of Fay and Herriot (1979). Many exploratory analyses of SAIPE data with alternative models of this type have been performed over the last decade (NRC Report 2000, Chapter 5), in order to choose the best available model specification and small-area predictors from the county-level variables derived from IRS and Food-Stamp data, and the most recent decennial census long-form data. There were four models which were judged best (NRC Report 2000, models (a)-(d), p. 56), from the point of view of maximized likelihood and adequacy of fit to the CPS response variables. The model adopted in SAIPE production uses as response variable the CPS weighted estimate of the number of poor related

school-age children in the county and log counts from IRS, Food-Stamp, and census records as predictors. Another model, which we focus on here, used as response variable Y_i in county i the ratio of the same CPS weighted estimate of number of poor children over the CPS-weighted estimator of the number of school-age related children, and as predictor a vector of variables \mathbf{X}_i with first component 1, and four other components

LTAXRT = logarithm of IRS-estimated child poverty rate;
LSTMPRT = logarithm of Food-Stamp participation rate;
LFILRT = logarithm of IRS child tax exemptions divided by estimated population <18;
LCPRT = logarithm of poverty rate for residents aged 5--17 from the latest decennial census.

The *SAIPE log-rate model* for estimates in income-years up to 1995 is

$$\log Y_i = x_i^{\text{tr}} \beta + u_i + e_i, \quad u_i \sim \mathcal{N}(0, \sigma_u^2), \quad e_i \sim \mathcal{N}(0, v_e/n_i) \quad (1)$$

in counties i in which at least one poor child was sampled. Here $\beta \in \mathbf{R}^p$ is a vector of unknown fixed-effect coefficients, and u_i, e_i are respectively PSU random effects and sampling errors, independent of each other within and across PSU's. Ordinarily, σ_u^2 is unknown and estimated, while v_e is known. (For SAIPE years after 1995, the denominator n_i in the variance of e_i has been replaced by $\sqrt{n_i}$.) In SAIPE, the parameter v_e is treated as unknown, and the model error variance estimated from fitting a regression model with the same predictors to the most recent previous decennial census data is treated as the known value of σ_u^2 .

We restrict attention here to the log-rate model (1) for three reasons.

(a) The SAIPE log-count model can be nested inside a model like (1) which augments the predictors X_i by the log denominators of the rates in (1). We address below (*cf.* Table 4 below) the comparative quality of fit to the data of various models — the log-count models, log-rate models, and others — which can all be nested within these augmented models.

(b) Rates should be more stable than counts across geography and time, since the frames for the census and CPS child poverty rates are different, but highly interpenetrating: differences in inclusion rules can result in differences in log-counts, but rates can still be similar within local areas.

(c) While counts can clearly not be stable over time in an expanding population, rates might be, and failing that, the relationship between the CPS and census rates might be.

Reasons (b)-(c) relate specifically to our goal of studying the stability of the relationships between CPS and census measurements of child-poverty, both across counties in and out of the CPS sample within model year, and across model years. Previously, Maiti and Slud (2002) restricted attention to the same log-rate model because rates are the natural small-area parameters to fit within a GLMM, and their objective was to compare SAIPE and GLMM small area estimates.

2. CPS vs SAIPE estimate vs Census

A central premise of the SAIPE approach is that the county child-poverty estimates which would be used if available are either those from the census, if performed in the current year, or from the CPS if the latter's sample could be greatly enhanced. In either case, whether the true child-poverty rate as measured by the census or by the CPS is considered to be the true parameter-value sought, the actual estimates are derived from CPS and administrative-records predictors and are checked against the CPS direct estimates and, in decennial years, against the census direct estimates. For this reason, among others, the relationship between census estimates and those derived from CPS and administrative-records predictors is of great interest. To the extent that this relationship remains relatively stable over time, in each intercensal year t the SAIPE small area estimates could be thought to predict the county log-rates that would have been measured in the full-scale census had it been conducted in year t . The stability of this census and SAIPE-predictor relationship can now be examined using the hard evidence available from the two decennial model years 1990 and 2000. Specifically, we check whether the SAIPE log-rate models are still adequate across the full range of counties to the same extent they were found to be in 1990, as documented in the NRC report (2000).

Various differences between 1990 census and CPS data on child poverty are documented under the SAIPE program web-pages (www.census.gov/hhes/www/saife/, specifically in the sub-page [techdoc/cencpsdf.html](http://www.census.gov/hhes/www/saife/techdoc/cencpsdf.html)), including differences based on data collection, poverty universe, and treatment of unrelated subfamily members. Tables given in those web-pages, and in 1997 internal Census Bureau papers of Robin Fisher and of Robert Fay, describe differences between census and CPS estimated counts, e.g. of poor and poor children. However, studies of corresponding differences between child poverty *rates* do not seem to be generally available.

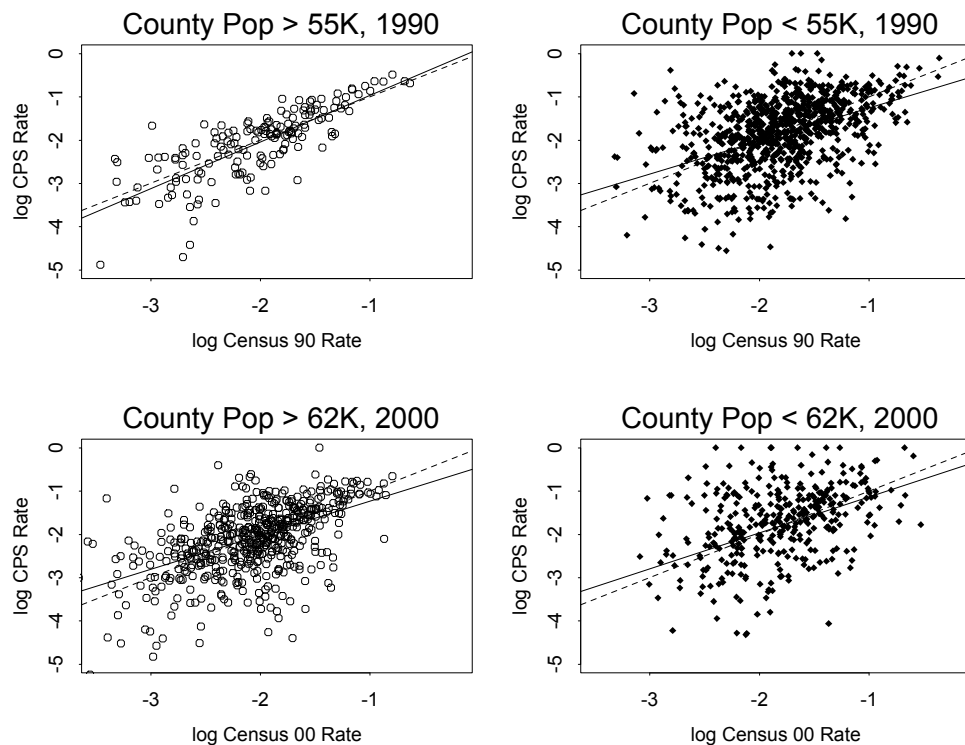


Figure 1: Scatterplot for 1990 and 2000 SAIPE-years, of log CPS child-poverty rates (i.e. $\log Y_i$ in (1)) versus log census rates adjusted to CPS poverty universe. Points are plotted only for counties with CPS-sampled poor children over years 1989-91 or 1999-2001. Separate plots are given for counties whose county census population was above (hollow circles) or below (solid diamonds) the upper quartile of counties (52000 in 1990, 60000 in 2000). In each plot, the 45° line is dashed, with the weighted least-squares line (weighted by CPS HU-sample size) solid.

In any given year, the non-CPS-sampled counties and the counties with no poor related school-age children in-sample are somewhat different from typical counties nationally: since the largest counties are always sampled, the non-sampled counties tend to be smaller and more rural, and that is also true of the sampled but dropped counties. (See the SAIPE and CPS web-pages saipe/techdoc/inputs/marcps.html and www.bls.census.gov/cps/cpsmain.htm for full details.) The SAIPE regression model, e.g. (1), is supposed to hold equally well over all counties, but is fitted using only those with sampled related poor school-age children. Moreover, although the relative modeling effectiveness on sampled versus non-sampled counties, in the decennial year 1990 when external comparisons on all counties were possible, was not a primary concern of the National Academy of Sciences (NAS) panel, the NRC Report (2000, p. 162) indicates that the standard SAIPE model over-predicts poverty in small counties.

The present research is concerned with data-analytic comparisons between the Census, CPS, and model-predicted child poverty-rate variables, over time and over categories of counties defined by size and CPS sample-inclusion.

3. SAIPE Data Analyses

We begin by direct scatter-plotting and linear modelling of the SAIPE log-rate responses $\log Y_i$ versus log census rates. The first result from such exploratory analysis is that while the 1990 log CPS rate is quite noisy and correlates only 0.52 with the 1990 log census rate, the least-square (or weighted least-square, weighted by the number of county CPS HUs) line is extremely close to the 45° line. That is, while the CPS data are noisy because many of the county samples are small, the best linear fit to log CPS rate in terms of log Census rate is close to the line equating the two rates. Figure 1 shows that this is true even when restricted to counties with 1990 population above or to those below the upper quartile, and also in 2000.

Next, Table 1 shows that the regression line is definitely closer to the line with intercept 0 and slope 1 when the regression is fitted by weighting squared residuals by the number of HUs in the county CPS sample. This weighting is precisely the one that would arise if the term u_i were absent from model (1), i.e., if σ_u^2 were replaced by 0. However, the value σ_u^2 treated as known in SAIPE, derived from estimating model errors in a model like (1) fitted to census data, is generally very small (about 0.014 in 1990, 0.016 in 2000). (This was also a key finding of Maiti and Slud 2002, that the SAIPE value of σ_u^2 in (1) is much smaller than would be estimated from internal data analysis of the model on CPS data alone.) Thus, the weighted linear regressions seem clearly the better choice.

We next consider linear models for $\log Y_i$ in terms of the SAIPE (log-rate) predictors \mathbf{X}_i , and guided by SAIPE methods and the previous two paragraphs, we fit weighted least squares models. These models are slightly simpler than (1), in that they do not contain a county random effect u_i . A first observation, clear from Table 2, is that both on 1990 and 2000 data, the weighted least-squares fitted coefficients are extremely close to the SAIPE log-rate model coefficients, while the ordinary least-squares coefficients are not. Our further objective is to examine the correspondence of these estimates with log census rates in 1990 and 2000 on several subsets of counties.

Published SAIPE model and small-area estimates are currently being calculated with sampling-error terms e_i in (1) assumed to have variances proportional not to $1/n_i$ but rather to $1/\sqrt{n_i}$, with supporting data-analyses given by Fisher and Asher (2000). The analyses of this paper, specifically those of Table 3, were repeated for that modified model and for the corresponding weighted linear regression models in which residuals-squared were weighted by $\sqrt{n_i}$ in place of n_i . Results are displayed only for model (1) as written, and for weighted regressions with weights n_i , because the slopes for the regressions weighted by $\sqrt{n_i}$ are similar but in almost every case differ from 1 (and the intercepts from 0) by more than the values shown, and in the same direction.

We do not study the SAIPE log child-poverty rate small area estimates directly in this paper (the EBLUP's of Ghosh and Rao 1994), even though those are (except for the slight difference between (1) and the log-count model actually used) the SAIPE program outputs directly used in Title I allocations. The reason is that the EBLUP's are by their definition convex combinations, with weights $\gamma_i = \sigma_u^2/(\sigma_u^2 + v_e/n_i)$ and $1 - \gamma_i$ in county i , between the log CPS rate estimates and their SAIPE-model predictors (where $\gamma_i = 0$ by definition in counties with no CPS sampled poor children). We have already indicated (in Table 1 and Figure 1) how the direct log CPS estimates behave with respect to logs of census rates, on larger and smaller CPS-sampled counties. We next study the relationship in years 1990 and 2000 between the predictors from model (1), and from the weighted least-squares predictors which would replace them if σ_u^2 were 0, versus the log census rates.

The predictors of the log CPS estimated child-poverty rates from the (weighted) linear regression are much less noisy than are the original $\log Y_i$ values. They also turn out to have very high correlations with the most current log census rates. These correlations (across all counties) are 0.91 in 1990 and 0.92 in 2000, when both the predictors and census are from the same model year. However, the predictors for SAIPE model years 1994 and 2000 respectively have correlations 0.956 and 0.942 with the log census rates from 1990, even though the correlation across all counties between the 1990 and 2000 census log rates is only 0.866. So it seems that the SAIPE-method log-rate predictors retain an unduly close relationship with the previous census, even toward the end of the decade.

Our most precise measure of the correspondence between *current* SAIPE log-rate model predictions and the true log child poverty rate is a joint plot of the model predictions and the log census rates in a decennial census year. Figure 2 gives such plots, both from 1990 and 2000 data, separately for counties with CPS-sampled poor children with population above and below the upper quartile of population sizes (55000 in 1990, 62000 in 2000). These plots show that in both years, in the group of larger CPS-sampled counties, there is nearly perfect linear agreement between model predictions and log census rates. (The dashed 45° line is virtually the same as the least-squares or weighted-least-squares line, and the error variance is rather small.) The linear agreement is also visually very good for the smaller counties, but close inspection of the Figure indicates that the slope of the linear relationship is strictly less than 1. Table 3 shows that the weighted least-squares lines expressing predictions from weighted linear regression model predictions or from the SAIPE log-rate model (1) have slopes within 2-standard-error tolerance of 1 (i.e. of the slope of the 45° line) for both groups of CPS-sampled counties in both 1990 and 2000.

Table 1. Coefficients for simple linear regressions of log CPS-rates on same-year log census rate. Lines fitted either by least-squares (*lin*) or least-squares with residuals squared weighted by number of HUs in CPS (*wtd*), separately for 1990 and 2000, and for counties with census resident population greater than (L) or less than (S) upper quartile of county populations.

Model	Intercept	Slope	SE(Slope)
90.L.lin	-.273	.901	.056
90.L.wtd	.108	1.078	.043
90.S.lin	-.636	.619	.068
90.S.wtd	-.498	.764	.065
00.L.lin	-.422	.816	.056
00.L.wtd	-.075	.979	.044
00.S.lin	-.551	.653	.083
00.S.wtd	-.443	.790	.080

Table 2. Coefficients for predictor variables in log CPS-rate regression models, in SAIPE data-years 1990 and 2000. Models were fitted by least-squares (*lin*), weighted least-squares with residuals squared weighted by numbers of HUs in CPS (*wt*), or Fay-Herriot model (*FH*) with variance σ_u^2 fixed = .014 in 1990 and = .016 in 2000.

Model	Intrept	ltaxrt	lstmprt	lfilrt	lcprt
90.lin	-0.13	0.39	0.21	-0.58	0.25
90.wt	0.22	0.27	0.29	-1.17	0.42
90.FH	0.18	0.29	0.29	-1.10	0.39
00.lin	-0.28	0.62	0.06	-0.03	0.22
00.wt	-0.09	0.49	0.12	-0.60	0.34
00.FH	-0.12	0.49	0.13	-0.47	0.33

Table 3: Coefficients for *predictors* of log CPS-rates from linear models with SAIPE log-rate predictors on same-year log census rate. Predictors were fitted by least-squares (*lin*) or weighted least-squares (*wtd*) or the SAIPE log-rate model (*saip*). The linear fits (and Standard Errors or SEs) for this Table were done via weighted least squares, separately for 1990 and 2000, for counties with CPS-sampled poor children with census population greater than (**Large**) or less than (**Small**) upper quartile, and also for non-CPS-sampled counties (**NonCPS**).

Model	Large			Small			NonCPS		
	lin	wtd	saip	lin	wtd	saip	lin	wtd	saip
1990									
Intercept	-.49	-.12	-.17	-.40	-.07	-.12	-.41	-.09	-.13
Slope	.79	.97	.95	.78	.96	.94	.79	.96	.94
SE(Slope)	.008	.010	.010	.013	.019	.019	.007	.010	.009
2000									
Intercept	-.45	-.23	-.28	-.50	-.30	-.35	-.51	-.31	-.35
Slope	.80	.91	.89	.72	.83	.82	.71	.82	.80
SE(Slope)	.008	.008	.008	.016	.019	.018	.006	.007	.007

Table 3 also shows that the slopes of the lines relating *unweighted* least-squares predictions of $\log Y_i$ in terms of SAIPE log-rate variables to log census rates are strictly less than 1.

Finally, Table 3 gives information about SAIPE log-rate model predictions versus log census rates in non-CPS-sampled counties. In 1990, the linear relationship is still virtually perfect, with slope close to 1. But in 2000, the corresponding slope is roughly 0.8, significantly less than 1, and the failure of the SAIPE log-rate model predictions to track log census rates closely can be seen in Figure 3, where the 45° line and the best fitting line through the scatterplot clearly differ.

We consider next a scatterplot which elaborates the difference already seen between fitted slopes in the **Large** and **Small** county portions of Table 3. Figure 4 displays the differences between SAIPE log-rate model predicted log child-poverty rates for 2000 and those observed in the 2000 census, plotted against the logarithm of county resident population. The dashed *lowess* curve gives a smoothed pieced-together version of the least-squares line fitted to the plotted points within narrow horizontal windows along the x-axis. This curve guides the eye in following an unmistakable nonlinear relationship between these log child-poverty rate differences and log county population. Points are plotted only for counties with CPS-sampled children, for ease of viewing, but the plot and *lowess* are very similar when points are plotted for all counties in 2000. Thus there is a systematic if small *nonlinear* dependence of the difference between log-rate model predictions and census log child-poverty rates on the logarithm of resident

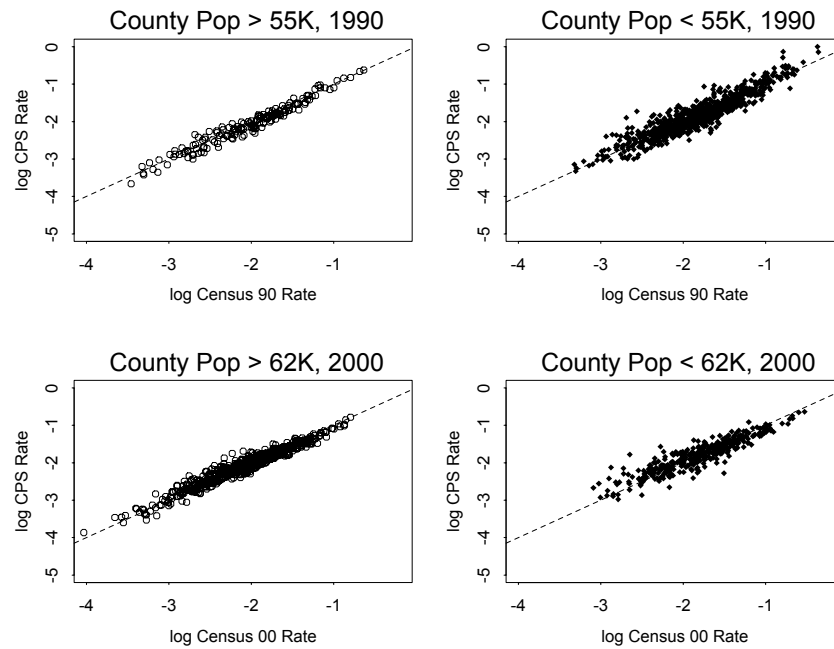


Figure 2: Scatterplot for 1990 and 2000 SAIPE-years, of weighted linear regression predictors for log CPS weighted child-poverty ratios (in terms of SAIPE log-rate model variables) versus log decennial census rates. As in Figure 1, points are plotted only for counties with CPS-sampled poor children, and are separated into groups according to year and whether county population is above or below upper quartile. Within each plot, the 45° line is plotted dashed.




Figure 3: Plot of weighted linear regression predictors of log CPS rates versus log Census rates. Predictors were fitted on SAIPE model year 2000 data, using only counties with CPS-sampled poor children. Plotted points are restricted to counties with no CPS sample in 1999-2001.

2000 Model Prediction Minus Census vs log Pop




Figure 4: Scatterplot of the difference between SAIPE model (1) predicted log child poverty rates in 2000 minus the log child poverty rates from 2000 census, versus logarithm of county resident population. Points are plotted only for CPS counties with sampled children, for ease of viewing. Vertical line corresponds to the logarithm of 62,000, the upper quartile of county population sizes. Dashed line is the `lowess` curve for the plotted points.

county population, in 2000. A very similar nonlinear dependence exists in the analogous picture for 1990, which is plotted as Figure 5.

In order to understand more clearly the nonlinear dependence we have just seen in Figures 4 and 5, consider the fits of several alternative Fay-Herriot models with auxiliary log population-size regressors on multi-year SAIPE data. We introduce three new models of log-rate type (i.e., with log-rate as response variable, like the SAIPE log-rate model), designated by prefix `Rt`, and one new model of log-Count type (response of log number of poor children), designated by prefix `Ct`. The SAIPE log-rate model (1) is designated `Rt.6`, and the SAIPE log-count model as `Ct.7`. The number in each model-name denotes the parameter dimension, including a sampling-error variance parameter v_e as in (1) but no u_i term. The new models, specified by their sets of regressors which are all taken from the denominators of terms in (1), are:

$$\text{Rt.8A} = \text{log-Rate terms} + \log(\text{Pop5:17}) + \text{LXMPT}$$

$$\text{Rt.8B} = \text{log-Rate terms} + \text{Lrespop} + \text{Lrespop}^2$$

$$\text{Rt.10} = \text{log-Rate terms} + \log(\text{CPSpop5:17}) + \log(\text{CNSpop5:17}) + \text{Lchpop} + \log(\text{CNSpoor5:17})$$

$$\text{Ct.10} = \text{log-Count terms} + \log(\text{Pop5:17}) + \log(\text{CPSpop5:17}) + \log(\text{CNSpop5:17})$$

All models were fitted as weighted linear regressions, without u_i terms. Both of the 10-parameter models `Rt.10`, `Ct.10` were chosen to contain both `Rt.6` and `Ct.7` nested within them. Table 4 below displays the negative-AIC quantities, equal to twice the maximized log-likelihood minus twice the parameter dimension, shifted by the constant 4600, for each of the multi-year SAIPE datasets. The comparison is intended to show both the comparability of the log-rate and log-count models *and* the helpfulness of using (in `Rt.8B`) linear and quadratic terms in log population size. It is not important to conclude here that `Rt.8B` is always the best model, since often the differences between

1990 Model Prediction Minus Census vs log Pop




Figure 5: Scatterplot of the difference between SAIPE model (1) predicted log child poverty rates in 1990 minus the log child poverty rates from 1990 census, versus logarithm of county resident population. Points are plotted only for CPS counties with sampled children, for ease of viewing. Vertical line corresponds to the logarithm of 55,000, the upper quartile of county population sizes. Dashed line is the `lowess` curve for the plotted points.

models are small, but only that in every year of data it is essentially as good as any other. In that way, the nonlinearity in Figures 4 and 5 is seen to be the dominant effect explained by the extra log population size terms of the augmented models.

4. Conclusions

The main conclusions of this research are: (i) that the SAIPE log child poverty rate model predictions and those of the SAIPE log-count model for CPS values are substantially the same, in all of the SAIPE model-years with the possible exception of 1994; (ii) that a weighted least squares fit to SAIPE data, departing from the coefficient estimators for model (1) only by substituting 0 for σ_u^2 , gives model-predictions of CPS log-rates essentially the same as does the model (1); (iii) that the SAIPE log-rate model continues in model years after 1994 to give predictions on CPS-sampled counties which track closely with 1990 census log rates; but (iv) that the SAIPE log-rate predictions on non-CPS-sampled counties may already not be tracking census 2000 log child-poverty rates closely enough in SAIPE model-year 2000; and (v) that there is evidence both in 1990 and 2000, of a slightly nonlinear systematic dependence of the difference between SAIPE log-rate model and Census log child poverty rates on the logarithm of county resident population size; (vi) which is borne out by a comparison of the fits, *in every year of SAIPE data*, of log-count and log-rate models augmented by the log-denominators of the regressor terms in model (1). Additional work with suitable loss-functions to measure prediction errors is needed to quantify the failure of fit suggested in Figures 3–5.

Table 4: Comparison of AIC values for alternative Fay-Herriot models including the *log-rate model* (1) (**Rt.6**), the analogous SAIPE *log-count model* (**Ct.7**), and the other models augmented by log population-size terms, displayed above in (**T**). The numbers in the Table are *maximized* $2*\log Lik - 2*(parameter\ dimension) + 4600$, and so are shifted by -4600 from the *negative* of the usual AIC. In this Table larger displayed values indicate “better” models, but comparisons are meaningful only within columns. Underlined entry is the highest in each column.

Model	90	94	96	98	99	00
Rt.6	559.30	302.77	935.29	1047.77	938.29	894.99
Rt.8A	583.17	375.84	955.08	1067.11	934.25	901.57
Rt.8B	584.80	<u>381.50</u>	<u>956.77</u>	<u>1067.51</u>	935.97	<u>902.83</u>
Rt.10	582.73	377.69	953.47	1066.41	<u>939.32</u>	898.83
Ct.7	413.89	50.89	595.96	748.08	624.52	626.35
Ct.10	<u>587.70</u>	380.89	955.32	1065.56	919.19	900.19

References

- Citro, C. and Kalton, G., eds. (2000) **Small-Area Estimates of School-Age Children in Poverty**, *Evaluation of Current Methodology* (National Research Council), Washington DC: Nat. Acad. Press.
- Fay, R. and Herriot, R. (1979) Estimates of income for small places: an application of James-Stein procedures to census data. *Jour. Amer. Statist. Assoc.* **74**, 341-53.
- Fisher, R. and Asher, J. (2000) Alternate CPS sampling variance structures for constrained and unconstrained county models. SAIPE Tech. Report, posted to SAIPE web-page www.census.gov/hhes/www/tecrep.html.
- Ghosh, M., and Rao, J.N.K. (1994) Small Area Estimation: an appraisal. *Statist. Sci.* **9**, 55-93.
- Maiti, T. and Slud, E. (2002) Comparison of Small Area Models in SAIPE. Census Bureau preprint.