Alternate CPS Sampling Variance Structures for Constrained and Unconstrained County Models: Additional Research

Jana Asher and Robin Fisher*

September 4, 2000

1 Introduction

In December of 1999 the Small Area Income and Poverty Estimates program (SAIPE) released a technical report titled "Alternate CPS Sampling Variance Structures for Constrained and Unconstrained County Models" (see Fisher and Asher, 2000) in which a preliminary proposal to change the SAIPE county model sample variance structure was created. The report recommended switching the variance function to the inverse of the square root of Current Population Survey (CPS) sample size, but also recommended further research before such a decision was made.

After releasing the technical report described above, SAIPE received commentary suggesting that further basic research was required. Specifically, closer examination of the effect on the regression coefficients and shrinkage weights of changing the sampling variance function was suggested, as well as an examination of the effect to estimates for counties with different CPS sample sizes. This document addresses that need for further basic research.

Section 2 of this document provides a summary of the original "Alternate CPS Sampling Variance Structures for Constrained and Unconstrained County Models" report and

*Jana Asher is a mathematical statistician in the Planning, Research, and Evaluation Division of the U.S. Census Bureau. Robin Fisher is a mathematical statistician in the Housing and Household Economic Statistics Division of the U.S. Census Bureau. This technical report is based on research and analysis undertaken by U.S. Census Bureau staff. It has undergone a more limited review than official Census Bureau publications. Research results and conclusions expressed are those of the authors and do not necessarily indicate concurrence by the Census Bureau. It is released to inform interested parties of current research and to encourage discussion. The authors wish to acknowledge the support and assistance of William R. Bell, Joseph Conklin, Richard Denby, Elizabeth Huang, Paul Siegel, Joanna Turner, and Daniel Weinberg.

explains the purpose of the additional research presented here. Section 3 contains a comparison of the regression coefficients and shrinkage weights of the alternative and original SAIPE models, a comparison of the original and alternative model results and Census estimates for the same year (income year 1989), and direct comparisons of the original and alternative models through relative differences and likelihood values. Section 4 contains the final recommendation to switch to the constrained model for the income year 1997 estimation procedure.

2 Background

As described in our original technical report, the current SAIPE model for county poor is:

$$\mathbf{y} = \mathbf{X}\beta + \mathbf{u} + \mathbf{e}$$

where $\mathbf{y}$ is a vector of the log of a three-year average of CPS number of poor for counties, $\mathbf{X}$ is a matrix of log values of variables from administrative records, $\mathbf{u}$ represents the model error and is distributed $\mathrm{N}[0, \mathbf{V_u}]$, and $\mathbf{e}$ represents the sampling error and is distributed $\mathrm{N}[0, \mathbf{V_{ec}}]$. $\mathbf{V_u}$ is assumed to take the form $v_u\mathbf{I}$; $\mathbf{V_{ec}}$ is assumed to be a diagonal matrix whose entries take the form $\frac{\sigma^2}{k_i}$ where k_i, the CPS sample size for county i, has the range (1, 4968). CPS sample size is simply taken as the sum of the sample sizes for the three years used in the estimation procedure. After regression estimates are formed, they are combined with the original CPS estimates via a weighted average called a "shrinkage estimate." The shrinkage estimates are then transformed back to original scale from the log scale; a bias correction term of the variance of the estimate divided by two is included in the transformation. Finally, the county level estimates are adjusted so that their sum conforms to state estimates; these final estimates are called the "raked estimates." The data used in this technical report are for related children aged 5-17 in families; the three-year average of CPS poor is centered around income year 1989. 1274 out of 3139 counties have CPS sample size greater than zero for income year 1989.

Initial analysis of the standardized residuals of this model suggested a problem; when the standardized residuals were plotted against the rank of sample size, a heteroscedastic pattern was revealed. The standardized residuals represent the sum of the model error and the sampling error divided by the joint standard deviation of this sum. Since the model error is assumed to have constant variance, we took the heteroscedasticity of the standardized residuals to suggest a misspecification of the sampling error variance.

As described in our original report, the current county level estimation procedure has two steps. The first is only for the estimation of the model error variance. In this step, we estimate the variance of census estimates of poverty with maximum likelihood where the 1990 decennial census direct estimate of poverty is the dependent variable and the sampling error variance is estimated with a generalized variance function. In the second stage, the model is estimated with CPS poor as the dependent variable and with the model error variance from the census model. The sampling error is assumed to be a constant times the inverse of the CPS sample size for the county; maximum likelihood estimation is used to determine an estimate of the value of this constant. We call the model created using this two-step procedure the constrained model, as the model error variance is constrained to a value created from census data. Another option is to assume that the model error is constant, but not estimate it through the census data and sampling error. To do this, maximum likelihood estimation is used to jointly determine v_u and σ^2. We call this model the unconstrained model.

The original report considered two alternate functions for CPS sampling variance: a piece-wise function of the estimated census sampling variance fitted to the residuals of the regression, and a constant times the inverse of the square root of CPS sample size. For the research described in the current document, we have dropped the piece-wise (Fuller/Goyeneche) function due to difficulty of implementation, and just consider the unconstrained and constrained modeling procedures with the inverse of the square root of CPS sample size as the basis for the sampling variance function.

Tests performed for the original technical report included a regression analysis of the squared standardized residuals on the rank of CPS sample size for each model. The test was based on the assumption that the squared standardized residuals follow a chi-squared distribution with one degree of freedom and used a Spearman's ρ test to determine if the slopes of these regression lines were significantly different than zero. Using this test, the slope for the original model was found to be significantly different than zero, while the slopes for the unconstrained and constrained model regressions were not. We also compared the results of the constrained and unconstrained model regression slopes and we failed to reject the hypothesis that $\beta_u = \beta_c$. Our conclusion was that the unconstrained and constrained modeling procedures produce homogeneous standardized residuals, while the original modeling procedure does not.

We also compared absolute relative differences between the alternate and original estimates in the original report, and concluded that the square root variance model led to about a 1% overall change in the shrunk estimates for both the constrained and unconstrained modeling procedure. We concluded that the sampling variance model had little effect on the value of the estimates. The results in Section 3 of the current document will dispute this conclusion.

As a final check in the original report, SAIPE created plots of the standardized, relative, and absolute relative differences of the estimates for a particular variance function from the 1990 census values, divided into categories such as census division, 1990 resident population, percent of population Black or Hispanic, percent in group quarters, and percent poor as given by the 1990 census. Results of these tests suggested that the alternate modeling options do not introduce bias into the county model; the plots for the different variance functions are quite similar for most demographic factors. We additionally noted that the proposed variance functions led to estimates that appeared to produce better results when comparing the standardized differences from the Census against 1990 resident population.

As a result of this research, we recommended further study of the unconstrained model taking sampling error to be proportional to the inverse of the square root of CPS county sample size because we believed that removing dependency on the census data would be beneficial. We concluded the first technical report with suggestions for future research, which included changing our current definition of CPS sample size to account for the sample overlap between years, testing the proposed variance models for different statistics and different estimate years (i.e., 1993 and 1995), testing alternate sampling variance models where the variance depends on the value estimated (e.g., binomial variance models), and using a generalized variance function to estimate sample variance.

The research that followed did not conform to our original intentions to further refine the unconstrained model. Following the suggestions of a colleague, we instead chose to further study the constrained and unconstrained models to determine how changing the sampling variance function specifically affected the estimates. The results of this new research follow in Section 3.

3 Results

A question of interest is the following: are the differences in the final estimates for the original, constrained, and unconstrained models due to changes in the regression coefficients (and therefore the underlying regression estimates), the shrinkage weights (and therefore the shrinkage estimates), or the bias correction term used in the transformation? Once this question is answered, the remaining question is which results are "better" - those from the original, constrained, or unconstrained model. The results presented below attempt to determine both the source of the changes between estimates produced by the three different modeling procedures, and the effect of these changes.

3.1 Regression Coefficients versus Shrinkage Weights

We first directly analyze the regression coefficients for the original, constrained, and unconstrained models using 1990 (income year 1989) CPS county data. Due to confidentiality issues, we cannot present these coefficients here, but we can describe their relationship. While the estimates of the regression coefficients for the constrained and unconstrained models are similar, they differ notably from those of the original model. The intercept for the original model is smaller than those for the alternative models, and, in all but one case, the coefficients are larger for the original model than for the constrained model (in one case, the coefficients are virtually the same for all three models). The smaller intercept and larger values for β for most of the variables in the original model mean that the original regression surface has a steeped slope through the cloud of data points than the alternative model regression surfaces. This will result in smaller population counties having higher predicted y values in the alternative models than in the original model. Further discussion of the differences in the regression coefficients for the alternate and original models can be found in Section 3.2.

For comparison purposes, we also plot the weights on the regression estimates used for shrinkage against CPS county sample size in Figure 1. As CPS sample size increases, there appears to be a growing difference between the three modeling procedures. Please note that the estimate for the model error variance for the unconstrained model is very close to zero, whereas the estimates for the model error variance for the original and constrained models is not; hence the weight on the model part of the shrinkage estimates for the unconstrained model must be close to one. The weights for the unconstrained models shown in the following plot indeed range between .998 and 1.

From the regression coefficients discussed above and the plot of the weights against CPS sample size shown in Figure 1, we tentatively conclude that differences in the final estimates between the three models presented are due both to the difference in the weights for shrinkage and the direct regression model results.

Figure 1 - Derived Weight on Model Part of Estimate, by Size of CPS County Sample, Income Year 1989.

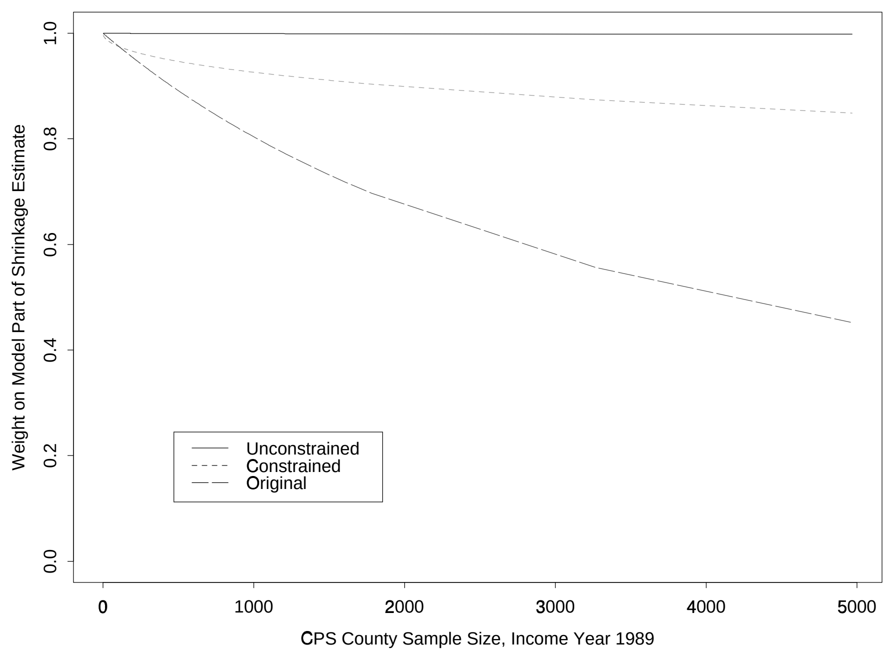

The differential change in the shrinkage weights over CPS sample size for the three models may suggest that counties with larger CPS sample size are more heavily affected by switching to one of the alternate models than counties with smaller CPS sample size. As a result, the remainder of the analyses presented here will include breakdowns by CPS sample size.

3.2 Comparison to Census Results

A question of interest that was not addressed in our original technical report is how the alternative and original models fare in a comparison to the census results for the same year. In this context we are using the census estimates of number poor as "truth" in spite of the fact that the census and CPS measure different definitions of poverty. This comparison is still useful for two reasons. The first is that the census estimates are as close to truth for small geographic areas as we have available. The second is that by comparing the original and alternative model estimates during different stages of the estimation procedure to the census, we are able to better determine how changing the sampling variance function affects the regression coefficients versus the shrinkage weights or bias corrections.

Figure 2 shows the relative differences from the census for each of the three models, plotted against log of CPS sample size. Log of CPS sample size is used for ease of plotting. The formula used for the relative difference of each county in the analysis is (model - census) / census. Note that the model estimates used are the raked estimates, and the census values used are adjusted to the CPS direct estimate of poverty at the national level. All counties with CPS sample size greater than zero are included.

The results of looking at the relative difference from the census on a point-by-point basis as shown below follow our expectations. The relative differences seem to be reasonably split below and above the zero line, but there seems to be a funnel-shaped pattern to the relative differences. Since the sampling variance of the estimates for small CPS sample size counties is modeled as higher than that of large CPS sample size counties, we would expect the differences between census and our estimates to be potentially bigger for the estimates with a bigger overall variance. As a result, these plots are not surprising, but it is difficult to draw any particular conclusions from these plots.

Figure 2 - Relative Difference of Model Estimate from Census Estimate, Income Year 1989.

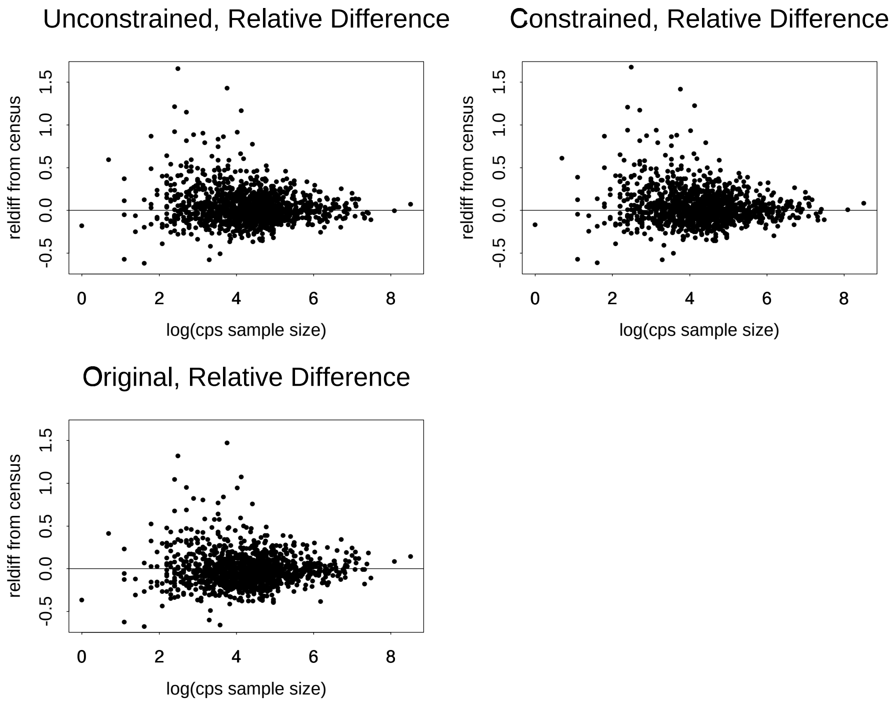

It is also of interest to compare summary statistics of the relative and absolute relative differences from the census as an additional check of the performance of the alternative and original models at several stages of the estimation process, for several sets of CPS sample size. A summary chart of this comparison is shown in Table 1. This table contains summary statistics of four sets of estimates: the estimates as given directly by the regression model after an exponential transformation, the estimates as given directly by the regression model after an exponential transformation that includes a bias correction term, the estimates after shrinkage and an exponential transformation that includes a bias correction term, and the estimates after raking. These estimates are compared to the census values adjusted to reflect the CPS poverty universe, as is done in the previous plots. The summary statistics are based on six sets of estimates defined by CPS sample size: all 3138 counties with estimates of census number poor greater than zero[1], the 1864 estimates with CPS sample size of zero, the 1274 estimates with CPS sample size greater than zero, the 637 of these estimates with sample size smaller than 80, the remaining 637 estimates with the largest CPS sample sizes, and the 16 estimates with CPS sample size greater than 1000. These divisions were picked somewhat arbitrarily to get a better sense of the effect of the change in the weights and regression coefficients on estimates with differing CPS sample sizes.

[1] Loving County, Texas has a census estimate of number poor of zero for 1989, and is therefore dropped from this table.

Table 1 - Mean of Absolute Relative Differences and Relative Differences of Income Year 1989 SAIPE Estimates and the 1990 Census by CPS Sample Size Category.

CPS Sample Size	Unconstrained[2]		Constrained[3]		Original[4]	
	Mean relative diff. from census	Mean abs. rel. diff. from census	Mean relative diff. from census	Mean abs. rel. diff. from census	Mean relative diff. from census	Mean abs. rel. diff. from census
Raked Estimates						
all	0.0587	0.169	0.0603	0.170	-0.00489	0.164
= 0	0.0714	0.190	0.0739	0.191	-0.00155	0.182
> 0	0.0402	0.138	0.0405	0.139	-0.00978	0.137
1-79	0.0626	0.170	0.0657	0.170	-0.00427	0.163
> 79	0.0178	0.107	0.0154	0.107	-0.0153	0.111
> 1000	0.0255	0.0783	0.0297	0.0811	0.0741	0.111
Estimates (After Shrinkage, before Raking)						
all	-0.0187	0.162	-0.00385	0.161	-0.0595	0.167
= 0	-0.0149	0.180	0.000592	0.176	-0.0659	0.184
> 0	-0.0241	0.139	-0.0103	0.138	-0.0500	0.142
1-79	0.00388	0.162	0.0205	0.163	-0.0425	0.167
> 79	-0.0522	0.116	-0.0411	0.112	-0.0574	0.116
> 1000	-0.0450	0.0971	-0.0211	0.0895	0.0794	0.113
Regression Estimates (Before Shrinkage and Raking, with bias correction)						
all	-0.0188	0.162	-0.0170	0.162	-0.0718	0.170
= 0	-0.0150	0.177	-0.0126	0.177	0.0783	0.187
> 0	-0.0243	0.139	-0.0233	0.139	-0.0624	0.144
1-79	0.00372	0.162	0.00587	0.162	-0.0557	0.170
> 79	-0.0523	0.116	-0.0525	0.117	-0.0691	0.119
> 1000	-0.0453	0.0973	-0.0508	0.101	0.0349	0.0820
Regression Estimates (Before Shrinkage and Raking, without bias correction)						
all	-0.0209	0.161	-0.0191	0.161	-0.0736	0.170
= 0	-0.0176	0.177	-0.0152	0.176	-0.0804	0.187
> 0	-0.0258	0.139	-0.0249	0.139	-0.0636	0.145
1-79	.00191	0.162	0.00404	0.162	-0.0572	0.170
> 79	-0.0536	0.117	-0.0538	0.117	-0.0701	0.119
> 1000	-0.0473	0.0989	-0.0529	0.103	0.0335	0.0816

[2] Mean relative difference from census, $\sum \frac{(\frac{unc-cen}{cen})}{n}$, and mean absolute relative difference from census, $\sum \frac{(\frac{abs(unc-cen)}{cen})}{n}$, where *unc* is the unconstrained estimate, *cen* is the census estimate for income year 1989, and n is the number of counties in the CPS sample size category.

[3] Equations are $\sum \frac{(\frac{con-cen}{cen})}{n}$, and $\sum \frac{(\frac{abs(con-cen)}{cen})}{n}$, where *con* is the constrained estimate.

[4] Equations are $\sum \frac{(\frac{org-cen}{cen})}{n}$, and $\sum \frac{(\frac{abs(org-cen)}{cen})}{n}$, where *org* is the original estimate.

There are several noteworthy trends in these statistics. The first is that, except for the last two groups of estimates of the largest CPS sample sizes, the mean of the absolute relative differences from the census is smaller for the original model than the alternative models for the final, raked estimates. The unconstrained and constrained absolute relative differences seem similar, however. Also for the raked estimates, the means of the relative differences are all positive for the alternative models, and larger in absolute value than those of the original model. It seems, after raking, that the alternative models fare worse, except for the counties with the largest sample sizes.

The results for the estimates after shrinking but before raking, however, tell a different story. In this case, the alternative models fare better than or equivalently to the original model in terms of mean of the absolute relative differences from the census, with the constrained model doing somewhat better than the unconstrained model. The relative differences also look better (closer to zero) for the alternative models. It is also interesting to note that the constrained model absolute relative differences before raking are comparable to or better than the absolute relative differences for the original model after raking.

For the estimates before shrinkage, the mean relative and absolute relative differences from the census are somewhat better in the unconstrained and constrained models than the original model. This suggests that the regression equation itself is not as good for the original model when compared to the constrained and unconstrained models, although these numbers are affected by the prediction variance for the regression equation as it is used for the bias correction factor. Examining the estimates created without the bias correction term shows almost identical results, confirming that the regression coefficient themselves are a large source of the difference between the original and alternate model estimates.

The conclusion is that while the original model estimates approach the census estimates at each processing step, for the most part the alternative model estimates move away from the census after raking (the exception being the group of counties with the largest CPS sample sizes). For this reason it is interesting to look at the raking factors to see if a cause for this trend can be elucidated. Some summary statistics for the raking factors are in Table 2.

Table 2 - Mean and Median Raking Factors for Adjusting to State SAIPE Estimates for Income Year 1989.

	Mean	Median
Unconstrained	1.080	1.090
Constrained	1.066	1.074
Original	1.061	1.041

The raking factors for all counties in a state are created by dividing the state estimate of poverty by the sum of the county estimates for poverty after shrinkage and the exponential transformation that includes the bias term. The mean and median values for these raking factors are higher overall for the unconstrained model, and slightly higher for the constrained model than for the original model. This suggests that the state estimate changes the estimates more in the alternative models than the original models by pulling the value of the estimates up. This trend is not uniform, however; for a few states, the estimates are reduced in size by the raking factors. Further analysis of the raking factors and their effect on the estimates would be informative.

A final comparison to the census can be performed as follows: a 90 percent confidence interval can be created for the difference between the estimates and the census values. The number of times that this confidence interval includes zero can then be calculated to determine a "coverage probability" for each estimation technique at different stages of estimation. Results of this analysis are in Table 3.

Table 3 - Percent that Cover Zero of the 90 Percent Confidence Intervals of the Difference Between Estimates and Census Values.[5]

	Percent Coverage		
Estimate	Unconstrained	Constrained	Original
After Raking	73.9	87.4	87.3
below zero	5.8	2.7	6.2
above zero	20.3	10.0	6.5
Before Raking	72.7	87.8	84.8
below zero	14.0	5.6	10.7
above zero	13.3	6.6	4.5

While the constrained and original models perform similarly, the unconstrained model performs poorly in comparison. It is interesting that the constrained model before raking performs comparably to the original model after raking.

3.3 Direct Comparison of Original and Alternative Models

As a final topic, we explore the relative differences between the original and alternative models. This topic was touched on briefly in our original report; using an aggregate absolute relative difference measurement, we concluded that switching from the original to the alternate model would result in about a 1% change in the estimates. This change, however, would be for the estimates on the log scale (before transformation). To determine

[5]The information in this table about the number of confidence intervals that fall above and below zero would suggest biases if we can assume that the cells in the table are independent. Since we cannot, this table does not suggest a consistent bias.

the effect on the estimates in their final state, we offer Table 4 as a more sophisticated and useful measurement. It contains the absolute relative differences and relative differences for the six different CPS sample size groups taken at the four estimate processing steps - before shrinkage without a bias correction, before shrinkage with a bias correction, after shrinkage with the bias correction, and after raking. All estimates used have undergone an exponential transformation.

Table 4 - Mean Absolute Relative Differences and Mean Relative Differences between the Original and Alternative Model Estimates.

	Unconstrained[6]		Constrained[7]	
CPS Sample Size	Mean relative diff. from original	Mean abs. rel. diff. from original	Mean relative diff. from original	Mean abs. rel. diff. from original
Raked Estimates				
all	0.0670	0.0751	0.0685	0.0762
< 1	0.0760	0.0801	0.0785	0.0824
> 0	0.0537	0.0678	0.0538	0.0672
1-79	0.0704	0.0759	0.0735	0.0794
> 79	0.0371	0.0597	0.0340	0.0549
> 1000	-0.0421	0.0564	-0.0391	0.0494
Estimates (After Shrinkage, before Raking)				
all	0.0469	0.0643	0.0628	0.0748
< 1	0.0585	0.0685	0.0754	0.0816
> 0	0.0301	0.0581	0.0446	0.0648
1-79	0.0514	0.0612	0.0689	0.0759
> 79	0.00884	0.0550	0.0202	0.0538
> 1000	-0.110	0.117	-0.0893	0.0979
Regression Estimates (Before Shrinkage and Raking, bias correction)				
all	0.0606	0.0715	0.0628	0.0746
< 1	0.0725	0.0782	0.0754	0.0816
> 0	0.0432	0.0617	0.0445	0.0644
1-79	0.0659	0.0721	0.0684	0.0751
> 79	0.0206	0.0513	0.0205	0.0537
> 1000	-0.0758	0.0863	-0.0810	0.0914
Regression Estimates (Before Shrinkage and Raking, no bias correction)				
all	0.0603	0.0714	0.0625	0.0745
< 1	0.0722	0.0780	0.0751	0.0814
> 0	0.0430	0.0616	0.0443	0.0643
1-79	0.0657	0.0720	0.0682	0.0750
> 79	0.0203	0.0513	0.0202	0.0536
> 1000	-0.0765	0.0868	-0.0818	0.0919

These summary statistics are similar for both the constrained and unconstrained model. Also, for every sample size group except the largest counties, the relative differences are all positive: the constrained and unconstrained models, overall, produce estimates that are larger than the original model. Since the total number of people represented by the raked estimates must be the same for each modeling procedure, the fact that the mean relative differences are positive and larger for the group of smaller counties suggests that the poor are being represented more in the smaller counties and less in the larger counties in the alternative models than in the original model. The change, overall, is about 7.5% for the final, raked estimates. Contrary to the findings in Fisher and Asher (2000), this represents a significant change.

For completeness, a plot of the relative and absolute relative differences between the two alternative models and the original model follows in Figure 3. All estimates used are the raked estimates; all counties with CPS sample size greater than zero are included in the plot. The relative difference in the estimates between the alternative models and the original model supports the concept of estimated poor being "sloshed" into small counties from large counties; basically it shows that the difference between the alternate and original estimates is positive for small CPS size counties, and negative for high CPS size counties. Since CPS sampling size and county population size are highly correlated, this suggests that taking square root transformation of the CPS sample size for the sampling variance function causes a reduction in size for estimates from high population counties, and an increase in size for estimates from low population counties, compared to keeping the unit CPS sample size.

---

[6]Mean relative difference from original, $\sum \frac{(\frac{unc-org}{org})}{n}$, and mean absolute relative difference from original, $\sum \frac{(\frac{abs(unc-org)}{org})}{n}$, where unc is unconstrained estimate, org is original estimate for income year 1989, and n is the number of counties in the CPS sample size category.

[7]Equations are $\sum \frac{(\frac{con-org}{org})}{n}$ and $\sum \frac{(\frac{abs(con-org)}{org})}{n}$, where con is the constrained estimate for income year 1989.

Figure 3 - Relative Differences between the Original and Alternate Model Estimates, After Raking.

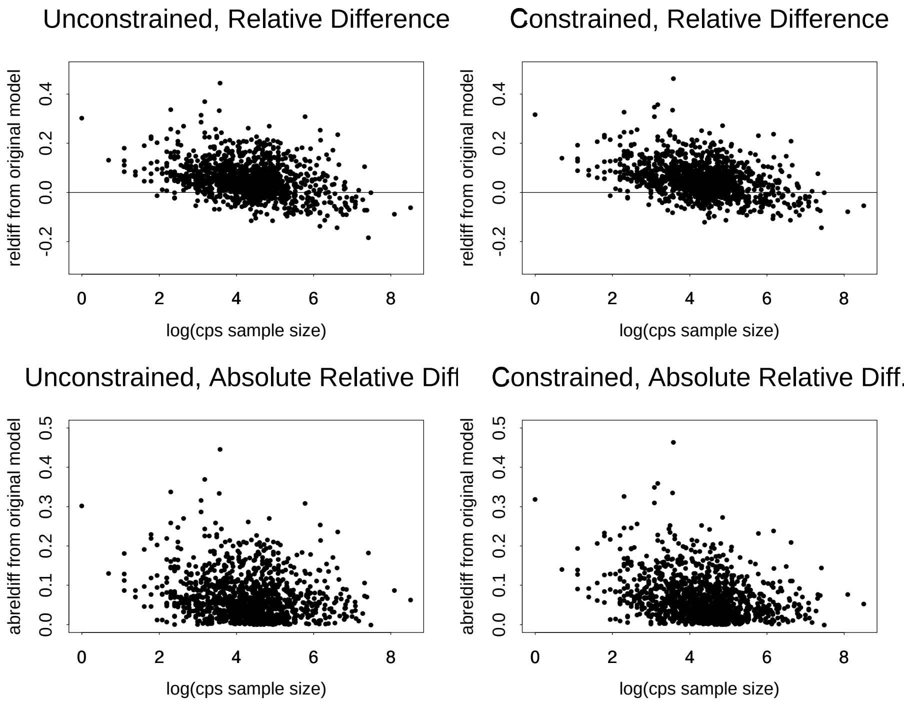

Finally, as part of a comparison of all three models, we note that the value of the likelihood at the maximum:

Unconstrained:	-1156.473
Constrained:	-1156.642
Original:	-1198.924

In terms of the likelihood, the constrained and unconstrained values are comparable, but both are considerably better than the value for the maximum of the likelihood function for the original model.

4 Recommendations

The work we have completed suggests that switching from the original model to one of the alternate models will cause a substantial change in the final estimates. This change will be differential: estimates for counties with smaller CPS county sample size (and therefore smaller overall population size) will be larger and estimates for counties with larger CPS county sample size (and therefore larger overall population size) will be smaller due to the change in the modeling procedure. The change will be due to both a change in the regression coefficients and a change in the variance structure. But the analyses performed both in the original technical report and in this document lead us to the following conclusion: *the assumption that the sampling variance can be modeled by the inverse of the square root of CPS sample size is better than the assumption that the sampling variance can be modeled by the inverse of the CPS sample size.* As a result, we recommend switching the sampling variance function to the inverse of the square root of CPS sample size for the 1997 income year estimate production.

At this time, we also recommend continuing to constrain the model using the census estimates, and believe that the unconstrained modeling procedure is inappropriate. We postulate that the unconstrained modeling procedure is untenable for a subset of the following three reasons. First, we believe that the model and sampling error terms in the unconstrained model may only be very weakly identified, causing the unconstrained model to poorly estimate the joint values of v_u and σ^2. By determining the value of v_u using the census sampling error estimates, the constrained and original modeling procedures avoid this issue. Second, the unconstrained modeling procedure may result in a marginal likelihood for v_u that is steeply unimodal close to zero, leading to unrealistically low estimates of the model error variance. A similar issue is seen in the current SAIPE state modeling procedure, and can be addressed by use of Bayesian estimation techniques (see Bell, 1999). Finally, there are assumptions in our modeling procedure that are imperfect, for example the assumption of normally distributed sampling error terms. We believe that the distribution of the residuals for the original SAIPE model may have large tails, which will affect the regression. In this case, constraining the model using the census values to determine the model error variance makes the modeling procedure more robust, so the outlying counties don't have as much effect on the estimation of the variance parameters. Further exploration of the joint and marginal likelihoods for the constrained and unconstrained modeling techniques would be required to determine the exact cause of the issues in the unconstrained modeling procedure.

Our conclusion is that the constrained model using the square root variance term provides the best modeling results before raking. It is important to note that the constrained model before raking appears to perform better than the original model after raking, both in the relative difference from the census statistics and in the mean percent coverage figures. As a result, we suggest a careful examination of the effect of raking on the accuracy and precision of the estimates for future income year estimation cycles.

5 Note

For completeness, we have reproduced the plots given in the original technical report on pp. 11-28 that contained diagnostic testing on demographic factors such as census division, 1990 resident population, percent of population Black or Hispanic, percent in group quarters, and percent poor as given by the 1990 census, but we divided the estimates into two groups. We divided the sample into counties where the CPS sample size (number of households) was from 0-79 households, versus CPS county sample size of 80 households or more. This resulted in the lower group having size 2502, and the upper group having 637 observations. The logic behind this split was that it divides the counties with CPS sample size greater than zero down the middle (there are 1274 counties with CPS sample size greater than zero for the 1989 data). Copies of these plots are available upon request; visual inspection does not provide evidence that the results from the different modeling procedures have a differential effect on counties with different CPS sample size.

6 References

Bell, William R. *Accounting for Uncertainty About Variances in Small Area Estimation.* Bulletin of the International Statistical Institute, 52nd Session, Helsinki, 1999.

Fisher, Robin. *Methods Used for Small Area Poverty and Income Estimation.* American Statistical Association, Proceedings of the Section on Government Statistics and Section on Social Statistics, 1997, pp 177-182.

Fisher, Robin and Jana Asher. *Alternate CPS Sampling Variance Structures for Constrained and Unconstrained County Models.* Technical Report #1, Small Area Income and Poverty Estimates program, 2000 (revised version).

Response to the NAB Request For the SAIPE Model Evaluations: Part 2. Part of mailing on August 27th, 1997, to the members of the National Academy of Sciences Panel on Estimates of Poverty for Small Geographic Areas, in preparation for the fifth meeting of the Panel held September 19-20, 1997.