Small Area Variance Modeling with Application to County Poverty Estimates from the American Community Survey

Jerry J. Maples, William R. Bell, and Elizabeth T. Huang*
Statistical Research Division, U.S. Census Bureau, Washington, DC 20233

Abstract

Variances in the American Community Survey are estimated using a replicate weight methodology (Fay, 1995). In counties with small sample sizes, the variance estimates of poverty statistics show wide variation as a function of sample size. Generalized Variance Functions (GVF) can be used to smooth out the uncertainty of the design-based variance estimate. We propose incorporating GVFs with small area model techniques to smooth out the variability in counties where the precision of the design-based variance is lacking. These smoothed variances can then be used in the small area models for poverty estimates.

keywords: Generalized Variance Functions, small area estimation

1. Introduction

The U.S. Census Bureau's Small Area Income and Poverty Estimates (SAIPE) program annually produces model-based estimates of income and poverty at the state and county levels for various age groups using Fay-Herriot (1979) models. Since 2005, data from the American Community Survey (ACS) have been used in the modeling. (Prior to 2005, data from the Current Population Survey were used.) For this paper, we will focus on county models used for estimating the number of related school-age (5-17 year old) children in poverty.

The county model used by the SAIPE program follows the Fay-Herriot framework:

$$\begin{aligned} y_i &= Y_i + e_i \\ Y_i &= X_i\beta + u_i \end{aligned} \tag{1}$$

where y_i is the log of the direct ACS estimate of the number of related children aged 5-17 in poverty (log total poor), Y_i is the log of the true number of related children aged 5-17 in poverty, $\mathrm{Var}(e_i) = \sigma_i^2$ is the sampling error variance in y_i as an estimate of Y_i, and $\mathrm{Var}(u_i) = \sigma_u^2$ is the model error variance. Usually, the sampling variances, σ_i^2, are treated as known, even though they are estimated. The estimates of the σ_i^2s are thus important inputs to the model.

In the ACS, sampling error variances are estimated using the successive difference replication variance estimator (Fay and Train, 1995). This method creates a set of replicate estimates, $y_{i,k}$, by perturbing the survey weights, according to values

*__Disclaimer:__ This report is released to inform interested parties of ongoing research and to encourage discussion. The views expressed on statistical, methodological, technical, or operational issues are those of the authors and not necessarily those of the U.S. Census Bureau.

in a Hadamard matrix, and then computing the variance estimate using the sum of squares of the replicate estimates around the original estimate. In application to ACS data, 80 replicate estimates are used. The resulting variance estimator is

$$\widehat{\text{Var}(y_i)} = \frac{4}{80}\sum_k (y_{i,k} - y_i)^2 \tag{2}$$

To obtain the variance of the log total poor, we apply the log transformation to the set of 80 replicate estimates before computing the sum of squares around the log of the original estimate.

$$s_i^2 = \text{Var}(\widehat{\log}(y_i)) = \frac{4}{80}\sum_k (\log(y_{i,k}) - \log(y_i))^2 \tag{3}$$

Figure 1 shows the estimated sampling error variances from (3) of the log number of children aged 5-17 in poverty as a function of sample size, defined as the number of responding households. Notice the large range of values of the estimated sampling variances for small and moderate sized counties (also note that both the X- and Y-axis have log scaling). Bell (2008) demonstrates the potential for a significant increase in the MSE (and a negative bias in the estimate of the MSE) of the small area estimate $E(Y_i|y_i)$ for counties whose sampling variances are substantially underestimated. This is important for counties in which the true variance is high (typically counties with small sample size).

One of the concerns with the Fay-Herriot model in (1) is that the estimated sampling error variance may be very imprecise for counties with small sample sizes. Although the ACS has a large national sample size, the sample sizes in individual counties for a single year can be small, with many (200+) counties having less than 50 responding households. (To address this issue, the official direct estimates of characteristics for these small counties from the ACS use 3 or 5 years of data, depending on the population size of the county.) Contrary to the usual assumption in small area modeling that the sampling variances are known, such small sample sizes raise questions about the precision of the resulting sampling variance estimates. Error in estimates of sampling error variances can adversely affect small area modeling in two ways. First, it can affect estimation of the regression coefficients β in (1), since these are fitted by weighted least squares, where the weights are the total variances, i.e. model error variance plus sampling error variance $\sigma_u^2 + \sigma_i^2$. Second, the sampling error variance is used to construct the weight, $w_i = \sigma_u^2/(\sigma_u^2+\sigma_i^2)$, given to the direct estimate when making the final shrinkage (empirical Bayes) estimate.

In this paper, we use a small area model framework to improve estimates of the sampling error variances of ACS county (log) poverty estimates. The model will incorporate a Generalized Variance Function (GVF) to explain the sampling error variances as a function of other variables. In the next sections, we will lay out the model framework, examine the issue of estimating the precision of the sampling variance estimator, and apply this model to the sampling variance estimates of log number of related children aged 5-17 in poverty from the 2005 American Community Survey.

2. Small Area Framework for Variance Models

Let s_i^2 be the estimated sampling variance of the log number of related children aged 5-17 in poverty for county i. We develop a model for the s_i^2 as follows. Given the true sampling variance, σ_i^2, we assume the variance estimate is unbiased, i.e., $E(s_i^2) = \sigma_i^2$, and that it follows a chi-squared distribution with d_i degrees of freedom. We will discuss what value to use for d_i in the next section. The second part of the model specifies the distribution of the σ_i^2's across counties. The σ_i^2's are assumed to follow an Inverse Gamma distribution centered around a GVF, $g_i = \exp(Z_i\gamma)$, and with precision parameter α. The predictors, Z_i, used in the GVF models will be discussed in Section 4.

The model can be written in hierarchical form as follows:

$$\begin{aligned} \frac{d_i s_i^2}{\sigma_i^2} | \sigma_i^2 &\sim \chi^2_{d_i} && (4) \\ \sigma_i^2 &\sim \text{InvGamma}(\alpha + 1, \alpha g_i). && (5) \end{aligned}$$

The quantity of interest, σ_i^2, is an unobserved random variable in this framework. To obtain the observed data distribution for s_i^2, we must integrate out σ_i^2 which has an Inverse Gamma$(\alpha + d_i/2 + 1, d_i s_i^2/2 + \alpha g_i)$ kernel. The variable transformation of $x = \frac{\alpha+1}{\alpha}\frac{s_i^2}{g_i}$ follows an F-distribution with degrees of freedom d_i and $2(\alpha + 1)$. The marginal distribution of s_i^2 then has the form:

$$f(s_i^2) = \frac{\Gamma(\alpha + d_i/2 + 1)}{\Gamma(\alpha + 1)\Gamma(d_i/2)} \frac{(\frac{d_i s_i^2}{\alpha g_i})^{d_i/2} \;\; s_i^{2-1}}{\left(1 + \frac{d_i s_i^2}{2\alpha g_i}\right)^{\alpha + d_i/2+1}} \tag{6}$$

The mean and variance of s_i^2 from this model framework are:

$$\begin{aligned} E(s_i^2) &= E(E(s_i^2|\sigma_i^2)) = E(\sigma_i^2) = g_i \\ Var(s_i^2) &= Var(E(s_i^2|\sigma_i^2)) + E(Var(s_i^2|\sigma_i^2)) \\ &= Var(\sigma_i^2) + E(2(\sigma^2)^2/d_i) \\ &= \frac{g_i^2}{(\alpha-1)} + \frac{2}{d_i}(g_i^2 + \frac{g_i^2}{(\alpha-1)}) \\ &= \frac{g_i^2(d_i + 2\alpha)}{d_i(\alpha - 1)} && (7) \end{aligned}$$

From (4) and (5), the conditional distribution of of $\sigma_i^2|s_i^2$ can be shown to be Inverse Gamma$(\alpha + d_i/2 + 1, d_i s_i^2/2 + \alpha g_i)$. This conditional distribution has expectation

$$E(\sigma_i^2|s_i^2) = \frac{d_i}{d_i + 2\alpha} s_i^2 + \frac{2\alpha}{d_i + 2\alpha} g_i \tag{8}$$

which is also the linear empirical Bayes estimate of σ_i^2 (using estimates for α and γ). Our goal is to estimate the parameters (γ, α) through maximum likelihood, so that we can generate estimates of σ_i^2 to replace the direct sampling variance estimates s_i^2.

Similar forms of this framework have been suggested by Otto and Bell (1995) and Arora and Lahiri (1997). Otto and Bell considered a multivariate version, modeling sampling covariance matrices for multiple years of state age-group poverty rate estimates from CPS. Their model was based on a Wishart distribution with an estimated constant degrees of freedom for all states. Arora and Lahiri used a chi-squared–inverse Gamma model in an application to data from the Consumer Expenditure Survey. They assumed the degrees of freedom to be $df_i = n_i - 1$.

The appropriateness of such simple assumptions about the degrees of freedom is questionable, as is suggested by simulation results in Huang and Bell (2009). For our application we wanted to develop more refined assumptions about the precision of the sampling variance estimator (i.e., its degrees of freedom). Thus, in the next section we investigate the precision of the sampling variance estimator for ACS estimates of log total poor using a bootstrap simulation. The results will guide us in making model assumptions that specify the degrees of freedom as a function of certain known quantities including sample size.

3. Bootstrap Simulation to Study the Precision of the Design-Based Sampling Variance Estimator

One problem in the framework described in Section 2 is deciding what to use for the degrees of freedom d_i for the variance estimator s_i^2? If the sample design was a simple random sample of size n, and the data were $i.i.d.$ $N(0, \sigma^2)$, then the variance of the total or mean would have $d_i = n - 1$ degrees of freedom. However, the ACS is not a simple random sample, but instead has a complex design, and the county point estimates we are modeling (of log total poor) are not just means of normally distributed observations. Some relevant features of the ACS sample design and estimation include differential sampling rates, systematic sampling, nonresponse follow-up subsampling, various weight adjustments, and raking to population controls (U.S. Census Bureau 2009). Additionally, we are estimating the variance of the log total poor which is a nonlinear transformation.

To investigate the precision of the ACS sampling variance estimator for county estimates of log total poor we did a bootstrap simulation study using 2005 ACS microdata. Our goal was to simulate those parts of the ACS design and estimation procedures that would be expected to have the largest effects on the variance estimator. We did not attempt to completely replicate every detail of the ACS sample design and estimation, something that would be extremely difficult. For our simulation study, we generated 1000 independent bootstrap datasets for each county and examined the precision of the sampling variance estimator across the simulations. The outcome variable of interest is the log number of related children aged 5-17 in poverty.

The first step in the bootstrap simulation for each county is to build the empirical distribution of households. Each household in sample has a number of related children aged 5-17 and a poverty status for that household. These are the only two variables we will concern ourselves with for this simulation. The weights for individuals within the same household are equal until the population controls are applied, which is one of the last weight adjustment procedures. For our simulation study, we will use the base weight (WSSF), after taking CAPI (computer assisted person interview) subsampling into account. To explain, for housing units that do not respond to the ACS by mail or telephone, a subsample is taken (for most areas a 1 in 3 subsample) and sent to CAPI follow-up. The weights on the CAPI cases are adjusted to reflect this subsampling procedure. The bootstrap sampling of households is done with replacement where the probability of selection for the household is proportional to the base sampling weight from the list of responding housing units.

For counties with small sample size, we were worried that the empirical distribution would not be a good approximation of the true population distribution of households. Therefore, for these counties we pool their data with data from the other counties that are within the same ACS Estimation Group. The ACS design had already grouped counties together (within the same state) by poverty rate, race and ethnicity distribution, distance from each other, etc., into estimation strata that are guaranteed to have a minimum of 400 person interviews. Additionally, these strata are the groups of counties in which the population controls are applied (rather than in each county individually). There were 2006 estimation groups in the 2005 ACS data for the 3141 counties. For bootstrap sampling we also conditioned on whether the data came from the subsampled CAPI cases, because the distribution of poverty is different between the CAPI and non-CAPI cases.

We draw a bootstrap sample for all counties in an estimation group using separate empirical distributions for the CAPI and non-CAPI households, and keeping fixed the number of CAPI and non-CAPI households. For the estimation group we compute the sample-based population estimate and use a constant ratio adjustment to all of the weights so that the sample-based population estimate equals the population estimate seen in the 2005 ACS. With this set of data and weights, we can then produce the design-based log total poor estimate and its estimated sampling variance using the successive difference replication variance method. We repeat this for 1000 samples. From these 1000 variance estimates, we can compute the mean, variance, and relative variance of the variance estimators. Since the assumed distribution of the variance estimator, given the true variance, is chi-squared with d_i degrees of freedom, and since the chi-squared distribution has a relative variance of $2/d_i$, we can approximate the degrees of freedom for the county from the simulation results. Figure 2 plots the estimated $\hat{d}_i$ as a function of sample size for the counties. Note the graph is on the log-log scale. This scale was chosen due to the strong linear association in addition to a nearly constant variance about the regression line.

A least-squares regression of the estimated degrees of freedom on sample size (on the log-log scale) indicates that the optimal power of sample size (where sample size equals number of responding households) is .43. We will use the square root of sample size as a reasonable approximation (since .5 is close to .43 and square root has an easier interpretation), and regress the estimated degrees of freedom on $\sqrt{n_i}$. This results in $\hat{df}_i = .36 \times \sqrt{n_i}$. We will assess how well this approximation works in the next section.

4. Variance Modeling for the ACS Log Poverty Estimates

The modeling of the variance is composed of two parts: selection of the predictors, Z_i, for the GVF and estimation of the parameters. Model fitting will be done by maximizing the observed data log likelihood based on (6). The dependent variable, s_i^2, is the estimated sampling variance of the log number of related children aged 5-17 in poverty from the 2005 ACS.

There are two main types of covariates that we consider: covariates that explain the mean number of poor (or poverty rate) and covariates that explain the survey design. The former covariates we take from the SAIPE production model of the log number of related children aged 5-17 in poverty:

1. **lfoodstamp**: log number of foodstamp participants
2. **lirschildp**: log number of IRS child exemptions in households in poverty
3. **lirschild**: log number of IRS child exemptions

4. **lrpop017**: log number of children aged 0-17 from demographic population estimates

5. **lcenpoor**: log number of related children aged 5-17 in poverty from Census 2000

The following are covariates based on the design and estimation procedures of the ACS:

1. **lhhct**: log number of responding households

2. **rrate**: percent of households in sample that gave a response (responding households / original sample size)

3. **capirate**: percent of responding households that are CAPI cases

4. **xpopcontrol**: percent of population the county contributes to its population control group

5. **persample**: aggregate sampling faction (of households) for county

A special variable called **lsaipemodelest** is the predicted log number of related children aged 5-17 in poverty from the 2005 SAIPE production model. This model uses the direct sampling variance estimates of the sampling error variances. It has the form:

$$\begin{aligned} \text{lsaipemodelest} &= -.421 + .173 \text{ lfoodstamp} + .548 \text{ lirschildp} - 1.037 \text{ lirschild} \\ &\quad +1.050 \text{ lrpop017} + .268 \text{ lcenpoor} \end{aligned} \tag{9}$$

Note that all of the predictor variables for **lsaipemodelest** are on a same scale as the population size of the county and thus are highly correlated with each other. A model-based estimate of poverty rate, **lsaipemodelrate**, uses the covariates listed above for the log number of children in poverty plus two more covariates (demographic population estimates for all ages and Census 2000 poverty universe for related 5-17 year olds) that become denominators for predictor rates. The modeled rate has the form:

$$\begin{aligned} \text{lsaipemodelrate} &= .372 + .556 \log(\text{Child Tax Poor Rate}) \\ &\quad +.167 \log(\text{Food stamp rate}) \\ &\quad -.414 \text{ Child Tax filing rate} \\ &\quad +.289 \log(\text{Census 2000 poverty rate}) \\ &\quad -.02 \log(\text{Pop 0-17}) \end{aligned} \tag{10}$$

In selecting predictors for the the GVF, we do not want to fit the error term in the direct sampling error variance. This can happen if a predictor, one of the variables in Z_i, is a survey response variable that has an error which is correlated with the error in the direct sampling error variance estimate. For example, rather than using the direct survey estimate, $\log(y_i)$, we can use the model regression estimate, *lregmodelest*. This issue also came up with another potential predictor, the number of households in poverty that contain at least one related child aged 5-17. This variable was more predictive in preliminary analyses of the estimated sampling variance than the number of responding households but we were concerned that there would be high correlation between the errors in the variance estimate and this version of 'sample size'.

To help with construction of the GVF, we look at Taylor's approximation of some known variance structures. If we assume the variance of the original (non-log) data is proportional to a power of the mean,

$$\begin{aligned} Var(y_i) &= k\mu^c \\ Var(\log(y_i)) &\approx k\mu^{c-2} = \exp(\gamma_0 + \gamma_1 \log(\mu)) \end{aligned} \tag{11}$$

where $\gamma_0 = \log(k)$ and $\gamma_1 = c - 2$. Allowing the γ's to be freely estimated yields model-based class of possible GVF's. The γ parameters will be estimated by maximizing the observed data distribution for s_i^2 given by (6). Suppose we instead assume a binomial type variance of the original data, then

$$\begin{aligned} Var(y_i) &= \frac{k\pi(1-\pi)}{n} \\ Var(\log(y_i)) &\approx \frac{k(1-\pi)}{n\pi} = \exp(\gamma_0 + \gamma_1 \log(\pi) + \gamma_2 \log(1-\pi) + \gamma_3 \log(n)) \end{aligned} \tag{12}$$

where $\gamma_0 = k$, $\gamma_1 = \gamma_3 = -1$ and $\gamma_2 = 1$ generalizes the binomial type variance. Variables such as sample size, response rate, percent of households in CAPI, and percent of households in sample can be viewed as predicting the design effect, k, the constant part of the variance function. For the variables that are rates, we can examine whether a transformation is needed (e.g. log rate) linear predictor $Z_i\gamma$ of the GVF.

Recall in the previous section, our bootstrap simulation suggested that the square root of sample size gave a reasonable approximation to the degrees of freedom, as computed from the empirical relative variance. In preliminary runs of the model, we noticed that the GVF model precision parameter α was unreasonable large thus putting very little weight on the direct estimate for the shrinkage estimator (median weight was 9%). One possible reason for this was that the simulation uniformly underestimated the precision of the direct variance estimator. Thus, we allow the degrees of freedom to be proportional, up to a scalar multiple, to the estimated degrees of freedom from the simulation or the approximation using the square root. This is now similar to fitting the fixed precision parameter for the direct variance estimator from Otto and Bell. In terms of parameters, we are adding a multiplicative scalar that must be estimated to the fixed degrees of freedom part of the model, i.e. instead of known d_i we now have κd_i. This factor will correct for any overall level of discrepancy in the variance which may be due to assumptions made in the simulation (systematic over or under estimation of the variance). Additionally, since the replication method uses 80 replicates, we limit the maximum value of the degrees of freedom to be 80 before evaluating the likelihood function (estimates of κd_i larger than 80 are set to 80 for likelihood function evaluation). This limit is imposed because the replicates ideally represent 80 orthogonal contrasts of the data, thus there are at most 80 'independent' sums of squares which translates to an upper bound of 80 for the degrees of freedom.

We will consider eighteen (3 x 3 x 2) variations of the GVF model for comparison based on three components. The first component is the GVF variance style: power-mean [mean], power-mean using the SAIPE predictors in (9) [mean-x] or binomial [bin]. The second component is the design variables: sample size [ss], sample size

Table 1: Log Likelihood, AIC and Parameter Estimates of GVF models

	Df = Sim				Df = sqrt n			
Model	loglike	AIC	α	scale	loglike	AIC	α	scale
mean: ss	2527.34	-2517.34	4.29	2.77	2540.44	-2530.44	4.26	2.88
ss+rates	2776.87[a]	-2758.87	5.65	2.92	2766.40	-2758.40	5.43	3.08
ss+log rates	2769.25	-2751.25	5.65	2.89	2768.28	-2750.28	5.44	3.05
mean-x: ss	2682.52[b]	-2664.52	5.21	2.78	2689.15[b]	-2671.15	5.29	2.80
ss+rates	2811.47[c]	-2785.47	6.05	2.85	2808.49[c]	-2782.49	5.83	3.00
ss+log rates	2792.83	-2766.83	5.91	2.85	2790.31	-2764.31	5.69	3.00
bin: ss	2511.41	-2487.41	4.06	2.89	2515.59	-2491.59	4.17	2.83
ss+rates	2637.38[d]	-2617.37	4.77	2.88	2641.00[d]	-2621.00	4.82	2.88
ss+log rates	2648.35[d]	-2628.35	4.79	2.99	2652.27[d]	-2632.27	4.85	2.91

[a] percapi not significant
[b] lrpop017 not significant
[c] percapi and lhhct not significant
[d] percapi and rrate not significant

and rates, or sample size and log rates. The final component is the degrees of freedom: proportional to the estimated degrees of freedom from the simulation [sim] or proportional to the square root of number of responding households in sample [sqrt n].

Table 1 compares the fits of the various models. A few major trends are apparent when comparing the likelihood and AIC values. First, variables (rates) related to the design and operations of the ACS significantly added explanatory power to the model. Second, the Binomial variance form did worse than using the mean or the variables that predict the mean. Third, using the rates in the GVF rather than the log rates consistently gave better fitting models. Finally, compared to the first two points, the differences between the likelihoods were small when switching the estimated degrees of freedom from the simulation result to the square root approximation.

The precision parameter α is highly correlated with the likelihood of the model ($\rho = .98$). The better models (higher likelihood) have higher precision / lower GVF model error variance (higher α). The scale parameters were consistent across the different GVF models. For the degrees of freedom estimated from the simulation, an inflation factor in the 2.8 to 3.0 range was suggested by the data. This means that the assumptions in the simulation ended up understating the precision of the survey variance estimator by a factor of almost 3.

The model [mean-x, ss + rates] was the best model GVF specification of the ones examined in Table 1. While the model using the square root of sample size did slightly worse, it is much simpler and does not require a large bootstrap simulation of a complex survey to obtain. Also, using the predictors of the SAIPE county model rather than the model prediction keeps the variance small area model uncoupled from the production small area model – not requiring simultaneous estimation or iterating between the SAIPE county and variance models.

5. Empirical Bayes Shrinkage of the Variance Estimates

After selecting our model [mean-x, ss + rates, Df = sqrt n], we can make the empirical Bayes estimates of $\hat{\sigma_i^2}|s^2$ from (8) using our parameter estimates for α and $g_i = \exp(Z_i\gamma)$. Since the number of replicates in the sampling variance estimator is 80, we limited the model-based estimate of degrees of freedom to 80. There are 40 counties with sample size larger than 5487 that reach this hard cap. Thus, the largest weight that will be given to a direct sampling variance estimate is 80/(80+2*5.83) = 87%. Figure 3 shows the distribution of weights on the direct variance estimator.

A comparison of the ratio of the empirical bayes estimate to the direct variance estimate is given in Figure 4. The smoothed variance estimates for some of the smaller counties show large underestimation of the true design variances by the direct variance estimator. This has two implications for the SAIPE county production model when treating the estimated variances as the true variance. First, the variances determine the relative weights given when fitting the model regression parameters. Second, the variances determine the weights that are given to the direct estimates versus the model-based estimates in the Empirical Bayes smoothing. Underestimating the variance for the small sample size counties could give them undue influence on the model fitting and pull the smoothed estimate too close to the direct estimate when the sample size perhaps does not warrant it.

6. Discussion

In this paper we have presented a small area model framework for estimates of sampling variance from complex surveys. Variance estimates for the number of related children aged 5-17 in poverty from the 2005 ACS were fitted using the model. We found that our approximation to the precision of the sampling variance estimator worked nearly as well as our degrees of freedom estimated from a bootstrap simulation of the ACS (see Table 1).

There is more research needed to understand the properties of the successive difference replication variance estimator, especially when applied to non-linear functions of the data. There were several limitation in the bootstrap simulation design. First, the simulation did not preserve the serial correlation in the sample list of housing units. Second, the weight adjustment procedure used in the simulation only incorporated a single cell population control, rather than the multiple weight adjustments including a multi-cell weight control. Despite the limitations of the simulation design, the bootstrap simulation of the ACS provided valuable insight as well as greater understanding of the different aspects of the design that could potentially influence the sampling variance. We plan to investigate the assumption of unbiasedness of the variance estimator for small sample sizes. Finally, we must understand the implication to the SAIPE county production model when replacing the direct variance estimates with the empirical Bayes estimates from the small area variance model.

References

Arora, V. and Lahiri, P. (1997) On the Superiority of the Bayesian Method over the BLUP in Small Area Estimation, *Statistica Sinica*, **7**, 1053-1063.

Bell, W.R. (2008), "Examining Sensitivity of Small Area Inferences to Uncertainty about Sampling Error Variances," *Proceedings of the Section on Survey Research Methods*, Denver, CO: American Statistical Association.

Fay, R.E. and Herriot, R.A. (1979) Estimation of Income from Small Places: An Application of James-Stein Procedures to Census Data, *Journal of the American Statistical Association*, **74**, 269-277.

Fay, R.E. and Train, G. (1995), "Aspects of Survey and Model-Based Postcensal Estimation of Income and Poverty Characteristics and States and Counties," *Proceedings of the Government Statistics Section*, Alexandria, VA: American Statistical Association, pp 154-159.

Huang, Elizabeth T. and Bell, William R. (2009), "A Simulation Study of the Distribution of Fay's Successive Difference Replication Variance Estimator," *Proceedings of the Survey Research Methods Section* [CD-ROM], Alexandria, VA: American Statistical Association, to appear.

Otto, M. and Bell W. (1995) "'Sampling Error Modeling of Poverty and Income Statistics for States," *Proceedings of the Government Statistics Section*, Alexandria, VA: American Statistical Association, 160-165.

U.S. Census Bureau, "Design and Methodology: American Community Survey," (2009), U.S. Government Printing Office, Washington, DC.

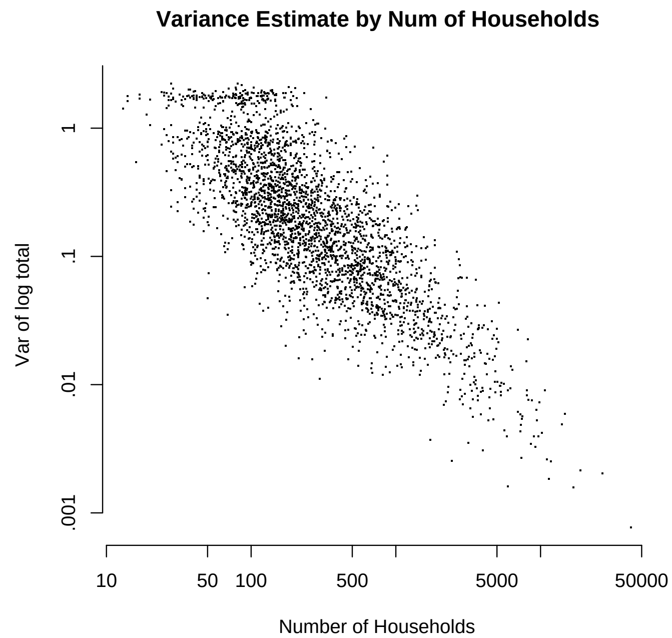

Figure 1: Sampling Variance Estimate by Number of Responding Households

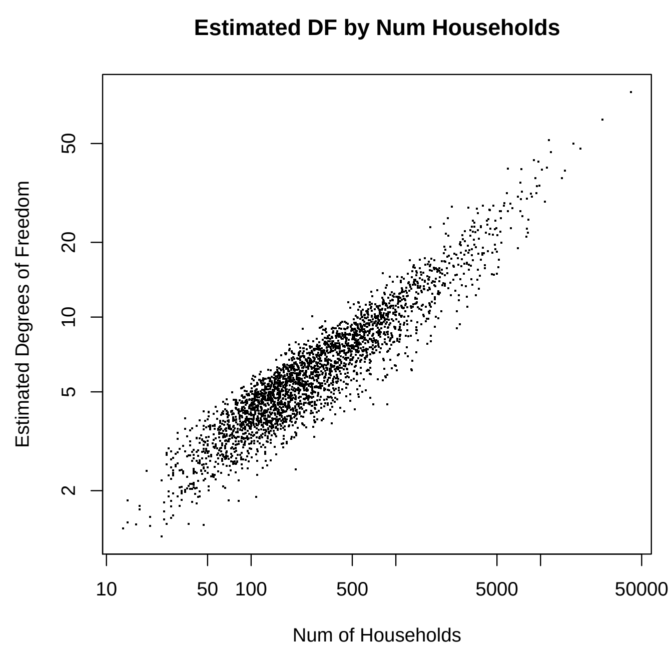

Figure 2: Estimated Degrees of Freedom by Number of Responding Households

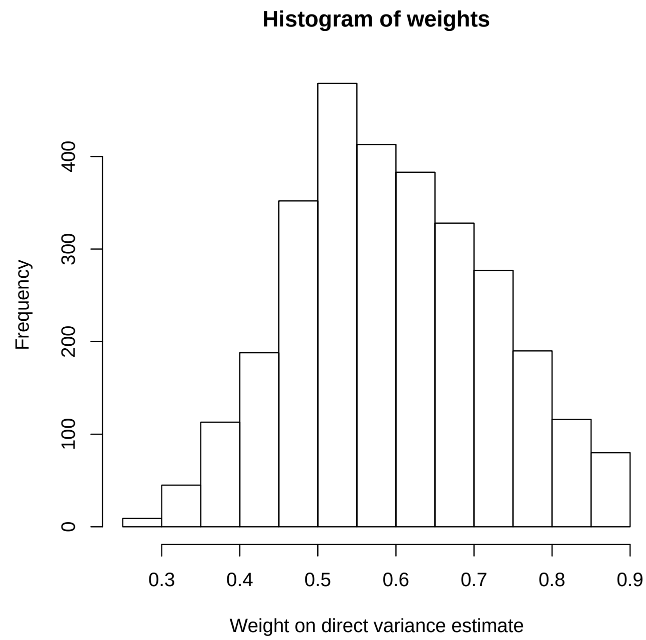

Figure 3: Histogram of Weights on Direct Estimate

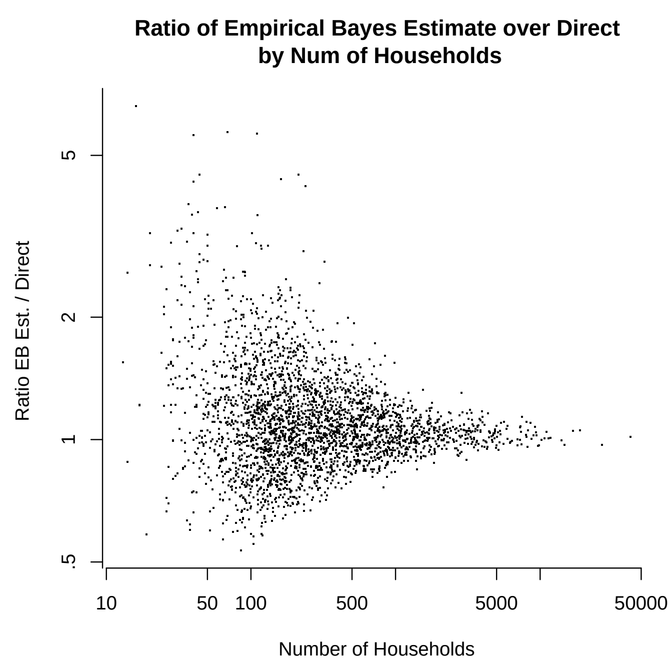

Figure 4: Ratio of Empirical Bayes to Direct Variance Estimates