RESEARCH REPORT SERIES
(Statistics #2011-02)

Using Small Area Modeling to Improve Design-Based Estimates of Variance for County Level Poverty Rate Estimates in the American Community Survey

Jerry J. Maples

Center for Statistical Research & Methodology
Research and Methodology Directorate
U.S. Census Bureau
Washington, D.C. 20233

Report Issued: March 2, 2011

Using Small Area Modeling to Improve Design-Based Estimates of Variance for County Level Poverty Rate Estimates in the American Community Survey

Jerry J. Maples

January 13, 2011

Abstract

In areas with small sample size, the estimates of sampling variability of rate estimates can have high uncertainty. Treating these estimated variances as the true sampling error variances in models for the underlying true rate can increase the mean squared error (Bell 2008). In counties with small sample sizes, the variance estimates of poverty statistics from the American Community Survey show wide variation even after accounting for sample size. Generalized Variance Functions (GVF) can be used to smooth out the uncertainty of the design-based variance estimates. We propose incorporating GVFs with small area model techniques to improve the variability of variance estimates in counties where the precision of the design-based variance is lacking. These smoothed variances can then be used in small area models for poverty rate estimates.

keywords: Generalized Variance Functions, small area estimation, poverty rates

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. Although the ACS has a large national sample size of about 3 million addresses, small area models are used to make estimates for the 3,141 counties, many of which do not have sufficient sample for adequate precision for estimates of poverty rates by various age groups. For this paper, we will focus on the estimates of poverty of related school-aged children (5-17 years old).

In the ACS, design-based sampling error variances are estimated using the successive difference replication variance estimator (Fay and Train, 1995). This method, in application to ACS data, creates a set of 80 replicate estimates, $y_{i,k}$, and then computes the variance estimate using the sum of squares of the replicate estimates around the original estimate.

$$\widehat{\mathrm{Var}(y_i)} = \frac{4}{80} \sum_k (y_{i,k} - y_i)^2 \tag{1}$$

Figure 1 plots the estimated sampling error variances of county poverty rates of children aged 5-17 as a function of sample size, defined as the number of responding ACS households. Notice the large range of values of the estimated design-based variance for small and moderate sized counties. Bell (2008) demonstrates the potential for a significant increase in the MSE (and a negative bias in the estimate of the MSE) of a small area estimate when the design-based variance is underestimated. This is important for counties in which the true variance is high (typically for counties with small sample size).

One of the concerns with the Fay-Herriot model is that the estimated sampling error variances 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 may be small with many (more than 200) counties having less than 50 responding households. In contrast, the official direct estimates of characteristics for these small counties from the ACS use three or five years of data, depending on the population size of the county. In small area models, it is standard practice to assume that the estimated sampling error variance is close enough to the true sampling error variance. Typically, estimating a second moment of the data (e.g. variance) is harder than estimating a first moment quantity such as a mean or a total.

Using the empirical Bayes estimates from a small area variance model for the design-based variances instead of the direct design-based error variance for the county model will have two implications in the small area point model. First, the regression coefficients in the small area model prediction are fitted using weighted least squares, where the weight is the total variance, i.e. the model error variance plus design-based error variance $\sigma_m^2 + \sigma_i^2$. Second, the design-based error variance is used to construct the weight, $w_i = \sigma_m^2/(\sigma_m^2 + \sigma_i^2)$, which is given to the direct estimate when making the final shrinkage (empirical Bayes) estimate.

In this paper, we will develop a small area model framework to improve estimates of the design-based sampling error variance for the poverty rate of school aged children. The model will incorporate a Generalized Variance Function (GVF) to explain the sampling error variances as a function of other variables.

2 Small Area Framework for Variance Models

Let s_i^2 be the estimated design-based variance of the poverty rate of related children aged 5-17 for county i, see (1). Given the true design-based variance, σ_i^2, we assume the estimated variance is unbiased, $E(s_i^2|\sigma_i^2) = \sigma_i^2$, and follows a Chi-Square distribution with d_i degrees of freedom. We will discuss the parameterization of d_i in the next section. The second part of the model describes the distribution of the σ_i^2s between the counties. The σ_i^2s are assumed to follow an Inverse Gamma distribution centered around the GVF, $g_i = g(X_i, \beta)$, with precision parameter α. The predictors, X_i, used in the GVF models will be discussed in Section 3.

$$\frac{d_i s_i^2}{\sigma_i^2} \sim \chi^2_{d_i} \tag{2}$$

$$\sigma_i^2 \sim \text{InvGamma}(\alpha + 1, \alpha g_i) \tag{3}$$

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 (Bell and Otto, 1995). 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

has the density:

$$f(s_i^2) = \frac{\Gamma(\alpha + d_i/2 + 1)}{\Gamma(\alpha+1)\Gamma(d_i/2)} \frac{(\frac{d_i s_i^2}{\alpha g_i})^{d_i/2} \;\; s_i^{2-1}}{\left(1 + \frac{d_i s_i^2}{2\alpha g_i}\right)^{\alpha + d_i/2+1}}. \quad (4)$$

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

$$\begin{aligned} E(s_i^2) &= E(E(s_i^2|\sigma_i^2)) = E(\sigma_i^2) = g_i \\ Var(s_i^2) &= Var(E(s_i^2|\sigma_i^2)) + E(Var(s_i^2|\sigma_i^2)) \\ &= Var(\sigma_i^2) + E(2(\sigma_i^2)^2/d_i) \\ &= \frac{g_i^2}{(\alpha-1)} + \frac{2}{d_i}(g_i^2 + \frac{g_i^2}{(\alpha-1)}) \\ &= \frac{g_i^2(d_i+2\alpha)}{d_i(\alpha-1)}. \end{aligned} \quad (5)$$

From the joint distribution created by (2) and (3), it can be shown that the the conditional distribution of $\sigma_i^2|s_i^2$ is Inverse Gamma$(\alpha+d_i/2+1, d_i s_i^2/2+\alpha g_i)$ (Otto and Bell, 1995). This conditional distribution has expectation

$$E(\sigma_i^2|s_i^2) = \frac{d_i}{d_i+2\alpha} s_i^2 + \frac{2\alpha}{d_i+2\alpha} g_i. \quad (6)$$

Substituting estimates for α and β and d_i into (7) gives an empirical Bayes estimate of σ_i^2. The empirical Bayes estimator is a weighted average between the direct variance estimate s_i^2 and the model-based GVF, g_i. The parameters α and d_i, which define the precison of the GVF and direct estimate respectively, control the weights in the weighted average. In fact, it is the relative values of these precisions which determine how much weight is given to the direct variance estimate compared to the modeled GVF. Since the degrees of freedom is increasing with sample size (Maples 2009) and α is constant for all sample sizes, counties with larger sample sizes will generally have larger weights on their direct estimates.

Similar forms of this framework have been suggested by Otto and Bell (1995) and Arora and Lahiri (1997). Other models to improve sampling variance estimates are discussed in Huang and Bell (2010). Otto and Bell considered a multivariate version, modeling multiple years of the variance for state level poverty rate estimates. Their model was based on a Wishart distribution with an estimated constant degrees of freedom for all states. This approach was considered appropriate because the difference in sample sizes between the states was not extremely large, and even the smallest state had a reasonable large sample size. Arora and Lahiri used a Chi-Square Gamma model and assumed their degrees of freedom to be $d_i = n_i - 1$, where n_i was the sample size. The American Community Survey has a different complex survey design than other surveys where this framework has been applied, e.g. the Current Population Survey (Otto and Bell, 1995) and the Consumer Expenditure Survey (Arora and Lahiri, 1997). It is not clear whether constant precision or precision analogous to simple random sampling applies to our model. Maples(2009) investigated the precision of the variance estimator for the log of a survey total in the ACS using a bootstrap simulation. The empirical results suggested that the effective degrees of freedom were proportional to the sample size to the .65 power and that the square root of sample size was a reasonable approximation. Huang and Bell (2009) have investigated the properties of Fay's variance estimator for rates.

3 Variance Modeling

The modeling of the variance is composed of two parts: selection of the predictors, X_i, for the GVF, and determination of the degrees of freedom for the design-based estimates, d_i. There are two main types of covariates that we consider: covariates that explain the poverty rate and covariates that explain the survey design. The former covariates we take from the SAIPE model (U.S. Census Bureau, 2010) for the poverty rate of related children aged 5-17:

1. **snaprt**: Supplemental Nutrition Assistance Program (SNAP) participation rate
2. **irschildpovrt**: percent of Internal Revenue Service (IRS) child exemptions that are in households in poverty
3. **taxfilert**: IRS filer rate for children
4. **rpop017**: number of children aged 0-17 from demographic population estimates
5. **cenpovrt**: poverty rate of related children aged 5-17 from Census 2000.

The following are covariates based on the ACS design and estimation procedures that affect the final weights of the sample cases in the ACS:

1. **hhct**: 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 fraction (of households) for county
6. **wratio**: ratio that measures the dispersion of base household weights, w_i, within a county,

$$\frac{\sum_i^H w_i^2}{\frac{1}{H}\left(\sum_i^H w_i\right)^2}.$$

In selecting predictors for the GVF, we do not want to explain sampling error variation in the direct estimate, s_i^2, of sampling error variance, σ_i^2. This can happen if a predictor, one of the variables in X_i, is a survey response variable that has an error which is correlated with the sampling error in the direct sampling error variance estimate, s_i^2. For example, one could use the number of households in poverty that contain at least one related child aged 5-17 as a measure of sample size. This variable was more predictive in preliminary analyses of the estimated design-based 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." Variables based on administrative record data or sources independent from the survey, e.g. IRS income tax data and demographic population estimates, should not have any correlation with error in the survey variance estimate and are fine to use in the GVF.

Since the true design-based variance is a function of the survey design and estimation procedures, we also wanted to include in the GVF covariates about the ACS operational procedures that affect weight adjustment. These include the percent of households collected in CAPI (CAPI subsampling weight adjustment), response rate (for nonresponse adjustment), dispersion of base household weights (the extent of unequal base weighting in an estimation area), sampling fraction, and the percent of population (based on demographic estimates) that the county contributes to its population control group. Variables that condition on the whether a household is a responding household are acceptable to use as long as the construction of the variables does not use the actual survey responses.

The GVF is assumed to have the form:

$$g(X_i, \beta) = p_i(1 - p_i)/n_i^* \tag{7}$$

where $p_i = p(X_{i1}\beta_1)$ is the underlying true poverty rate and $n_i^* = n^*(X_{i2}\beta_2)$ is the effective sample size. The function $p(\cdot)$ is the inverse logistic function and X_{i1} ={intercept, logit(snaprt), logit(irschildpovrt), log(taxfilert), logit(cenpovrt), log(rpop017)}. These are the variables that are under consideration for the poverty rate model. The reason for the log transformation instead of the logit for the tax filer rate is that nearly 10 percent of the counties have more child tax exemptions than the demographic population estimate for ages 0 to 17. This is to be expected because persons over the age of 17 can be included as dependent children, thus child exemptions, on the tax forms. The function $n^*(\cdot)$ is the exponential link, exp(), and X_{i2} = {intercept, log(hhct), [log(hhct)]^2, log(rrate), log(capirate), log(xpopcontrol), log(persample), log(wratio)}. These variables are to predict the effective sample size for a given county. This particular structure for the GVF has some advantages. First, it is a familiar form for the variance of rates. Second, it allows separation of the rate, which is a population quantity and so does not depend on the survey design, from the effective sample size, which solely depends on the survey. Third, by explicitly making the GVF a function of the true poverty rate, it provides a natural mechanism to extend the variance model to a joint model for both the poverty rate and its design-based variance.

Simulation studies by Huang and Bell (2009) show that the degrees of freedom is related to underlying true poverty rate. Here, the degrees of freedom, d_i is assumed to have the form:

$$d_i = \min(\gamma p_i \times \mathrm{hhct}^{\kappa}, 80) \tag{8}$$

where γ is the proportionality constant and κ is the power of the sample size. From the simulation study in Maples (2009), $\kappa = .6$ to $.7$ gave the optimal fit and the square root power ($\kappa = .5$) also gave a reasonable fit. Since the variance estimator is based on 80 replicates, the degrees of freedom, i.e. relative precision, is limited to 80.

3.1 Model Fitting

Model fitting will be done by maximizing the observed data log likelihood of $\sum_i \log(f(s_i^2))$ where $f(s_i^2)$ is from (4). The dependent variable, s_i^2, is the estimated sampling variance for the poverty rate of related children aged 5-17 in poverty from the 2005 ACS. Of the 3141 counties, 2973 counties had a design-based variance estimate of the poverty rate that was not zero, i.e., the estimated poverty rate was not 0% or 100%.

The profile likelihood for κ was computed to determine the "optimal" value. Figure 2 shows a graph of the profile likelihood as a function of κ. The optimal value of κ is around .62 (minus log likelihood -13833.77) and twice the difference in log likelihood from $\kappa = .5$ (can be considered a prior assumption) is 8.04, which is a statistically significant difference based on the likelihood ratio test with 1 degree of freedom. Currently, the number of responding households raised to a power is an empirical result and is not justified by any statistical theory based on the nature of the replicate weight variance estimator. Theoretical results have been worked out for the degrees of freedom in Normally distributed in Huang and Bell (2009).

Variable selection (backward selection) from the list of candidate covariates mentioned in the beginning of the section is done using the likelihood ratio test. Table 1 contains the log likelihood fits. The full model with all of the covariates is the starting model. A value of less than 3.84 for twice the difference between the full model likelihood and the model with a candidate variable removed, column -2ΔL, indicates that the variable did not improve the model given the other covariates. Three covariates were removed that meet this criteria: log(snaprt), log(rpop017) and the intercept term from $n^*(\cdot)$. It is not surprising that log(rpop017), a measure of size, is not predictive of the poverty rate, but it was tested to ensure that there was not a systematic trend with respect to county population size. The reduced model (R) with the three covariates removed

Table 1: Variable selection for the GVF

	Variable	Log L	-2ΔL	Log L(R)	-2ΔL (R)
	Full Model	13833.76			
$p(\cdot)$	intercept	13826.06	15.40	13817.80	29.30
	logit(snaprt)	13832.99	**1.44**		
	logit(irschildpovrt)	13803.37	60.78	13787.58	89.80
	log(taxfilert)	13830.59	6.78	13829.64	5.68
	log(rpop017)	13833.60	**0.32**		
	logit(cenpovrt)	13822.85	21.82	13819.07	26.82
$n^*(\cdot)$	intercept	13833.31	**0.90**		
	log(hhct)	13826.73	14.06	13769.20	126.56
	$[\log(\text{hhct})]^2$	13823.90	19.72	13795.95	73.06
	log(rrate)	13831.44	4.64	13827.61	9.74
	log(capirate)	13827.90	11.72	13824.41	16.14
	log(xpopcontrol)	13813.80	39.92	13803.20	58.56
	log(persample)	13829.74	8.04	13827.90	9.16
	log(wratio)	13826.70	14.12	13824.59	15.78
	Reduced Model (R)	13832.48	2.58		

has a twice difference of log likelihood of 2.56, which is less than 7.81 and thus meets the AIC criterion. The second set of columns, Log L(R) and -2ΔL(R) show twice the difference in log likelihood from the reduced model (R) when removing a single covariate. None of the remaining covariates meet the criteria to be removed. The variable selection results are identical when using $\kappa = .5$ instead of $\kappa = .62$.

The parameter estimates of the GVF for $p(\cdot)$ and $n^*(\cdot)$in the reduced model (R), after removing the nonpredictive covariates, are given in Table 2. The direction of the parameters for $p(\cdot)$ are as expected. The sign of the parameter estimates for $n^*(\cdot)$ are interesting as they relate characteristics about the survey and its operation to effective sample size. The rate of responding households and the more a county is not influenced by other counties for population controls increases the effective sample size. Figure 3 plots the log of the effective sample size against the log of the responding number of households. The log-log linear regression gives a fit of

$$n_i^* \approx \frac{\text{hhct}^{.941}}{5},$$

which could be used as a approximate estimate of n_i^* based on one predictor rather than the full set of variables based on survey operations. Figure 4 plots the estimated poverty rate against the direct ACS survey estimate of the poverty rate. Note that this estimate of p_i is based only on the variance estimates, s_i^2. Using estimates from a small area model of the poverty rate for p_i or doing a joint model of both the poverty rate and its design based variance should improve the estimation of p_i.

Figure 5 compares the model fit of the GVF to the design-based variance estimate. The model appears to fit through the center of the point mass with some counties having unusually small design-based variances. The GVF model smoothes out these unusually small variance estimates, where one may believe that the design-based estimate is under-estimating the true sampling error variance. Figure 6 shows the ratio of direct variance estimate to the GVF as a function of sample size. Figure 7 shows this plot for the largest 49 counties with sample size above 5000. Notice that all of the GVF's for the large counties are within a factor of two of their design-based variance estimates.

The estimated degrees of freedom for the direct variance estimates are shown in Figure 8 as a function of sample size. The degrees of freedom stay fairly low until sample size of 2000 and only 8 percent of counties have an estimated degrees of freedom larger than 15. Only seven counties hit the maximum of 80 degrees of freedom. Figure 9 shows the percent of weight on the direct variance estimate for the empirical Bayes variance estimate as a function of sample size and Figure 10 shows the distribution of the smoothing weights. With $\alpha = 14.26$, the largest weight on a direct variance estimate is 74 percent when doing the empirical Bayes smoothing. However, with most of the degrees of freedom less than 15, the majority of weights are less than 35 percent. These weights on the direct variance estimate are much smaller than is typically seen in SAIPE's small

Table 2: Parameter estimates for the reduced model (R)

	Variable	Parameter	Std. Err.
$p()$	intercept	.80	.12
	logit(irschildpovrt)	.88	.09
	log(taxfilert)	-.76	.31
	logit(cenpovrt)	.38	.07
$n^*()$	log(hhct)	.48	.04
	$[\log(\text{hhct})]^2$	.03	.003
	log(rrate)	-.47	.15
	log(capirate)	-.17	.04
	log(xpopcontrol)	.19	.02
	log(persample)	-.13	.04
	log(wratio)	-.36	.09
precision	α	14.26	2.01
	γ	.74	.05

area models for the point estimates for either log-level or log-rate. One explanation for the smaller smoothing weights on the design-based variance estimates is that it takes much larger sample size to get precision of the sampling error variance than design-based estimates of rates or totals. It could also be an indication that the assumption about the precision of the direct variance estimate is wrong. Figure 11 plots the empirical Bayes variance estimates as a function of sample size. Compared to Figure 1, it appears that the empirical Bayes smoothing with the GVF moved almost all of the unusually small direct variance estimates from counties with sample size between 100 and 1000 closer to the rest of the data cloud. Figure 12 compares the empirical Bayes variance estimates to the direct variance estimates. The empirical Bayes variance estimates of the largest 49 counties are all within a magnitude of 1.5 of the direct variance estimate (Figure 13).

4 Conclusions

In this paper we have presented a small area model framework for estimates of sampling variance for poverty rates from the American Community Survey. A generalized variance function (GVF) was developed specifically for the domain of related children 5 to 17 years old. Variables that corresponded to survey weight adjustment procedures were found to be significant predictors of the effective sample size. We found that the relative precision of the design-based variances was related to the responding sample size to the .62 power. The empirical Bayes estimates smooth the design-based estimates towards the GVF. The explicit parameterization of the underlying true poverty rate p_i will allow easy integration of the poverty rate model to form a joint model. Or one could use estimates of p_i from a small area model and iterate between the poverty rate model and variance model. Either of these two methods should give better prediction of the p_i's rather than estimating them solely from the model for the variance data.

One may expect the GVF to better fit the design-based variance estimates of the largest counties, and one may also expect the smoothing weights on the direct variance estimate to be larger for the largest counties. Bell (2008) has shown that incorrectly estimating the design-based variance by a factor of no more than two, only increases the mean squared error of the small area estimate (e.g. the poverty rate) by at most ten percent. The main purpose of this model is to correct the large under and over estimates of the design based variance likely to occur for counties with small sample sizes. These are the poor variance estimates that can greatly increase the mean squared error for the small area estimate of the poverty rate. The parameter of primary concern is the underlying poverty rate (from a separate model, not the rate, $p(\cdot)$ estimated within this model), not the design-based variance; thus, we only need the model of the design-based variances to be a reasonable representation.

The specification of the degrees of freedom, especially the κ parameter needs to be studied further. The first step will be to fit the model to additional years of ACS data and see if the data supports a constant κ over all of the years. There may also be some semi- and non-parametric methods that can be used to link the relative precision with sample size.

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, E.T. and Bell, W.R. (2009). "A Simulation Study of the Distribution of Fay's Successive Difference Replication Variance Estimator," *Proceedings of the American Statistical Associ-*

ation, Survey Research Methods Section, [CD-ROM], Alexandria, VA: American Statistical Association

Huang, E.T. and Bell, W.R. (2010). "Further Simulation Results on the Dstirubtion of Some Survey Variance Estimators," *Proceedings of the American Statistical Association, Survey Research Methods Section*, [CD-ROM], Alexandria, VA: American Statistical Association

Maples, J.J., Bell W.R. and Huang E.T. (2009). "Small Area Variance Modeling with Application to County Poverty Estimates from the American Community Survey," *Proceedings of the American Statistical Association, Survey Research Methods Section*, [CD-ROM], Alexandria, VA: American Statistical Association

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.

U.S. Census Bureau. "Small Area Income and Poverty Estimates." (2010). http://www.census.gov/hhes/www/saipe, accessed on Sep 27, 2010.

Figure 1: Design-based Variance Estimate by Number of Responding Households

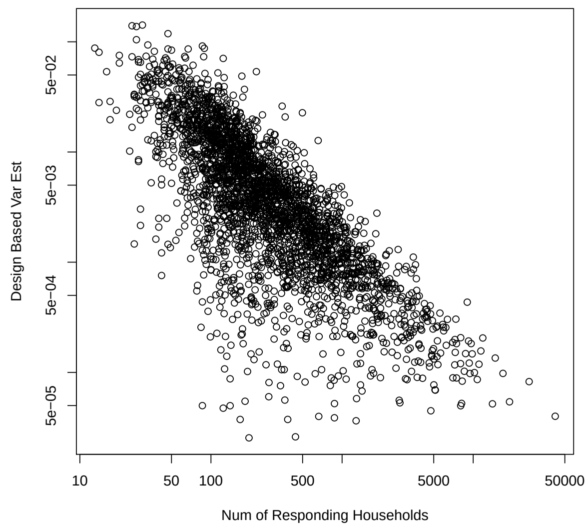

Figure 2: Profile Minus Log Likelihood of κ

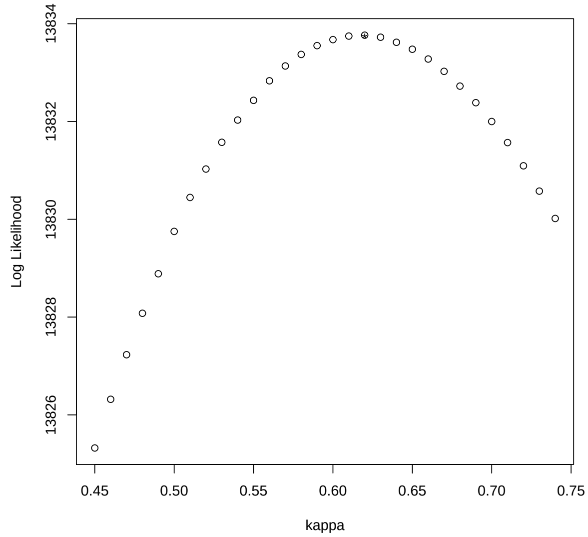

Figure 3: Estimated Effective Sample Size vs Responding Sample Size

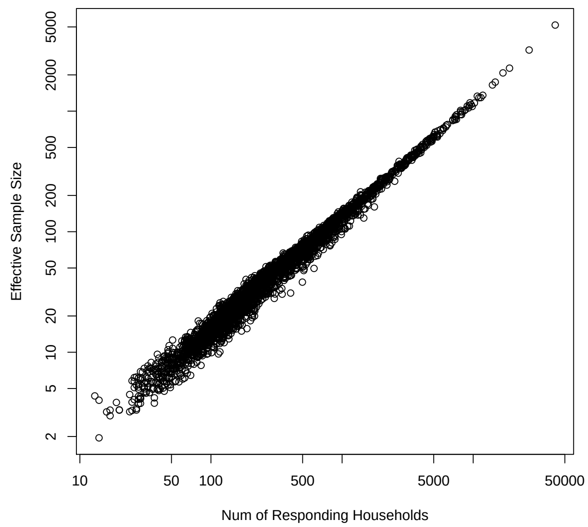

Figure 4: GVF Estimate of Poverty Rate Compared to Direct

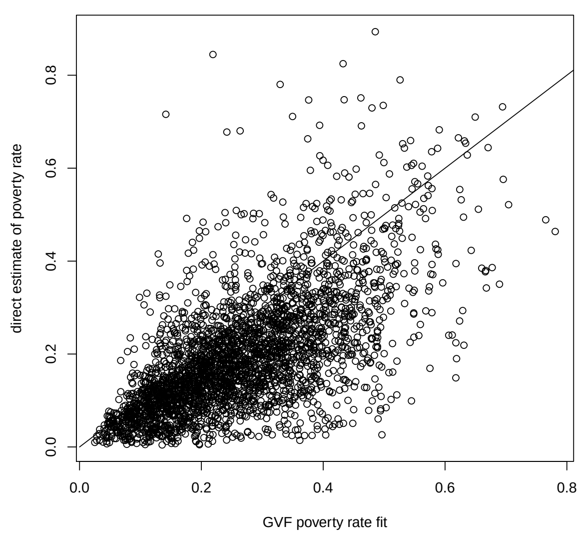

Figure 5: GVF Fit Compared to Direct

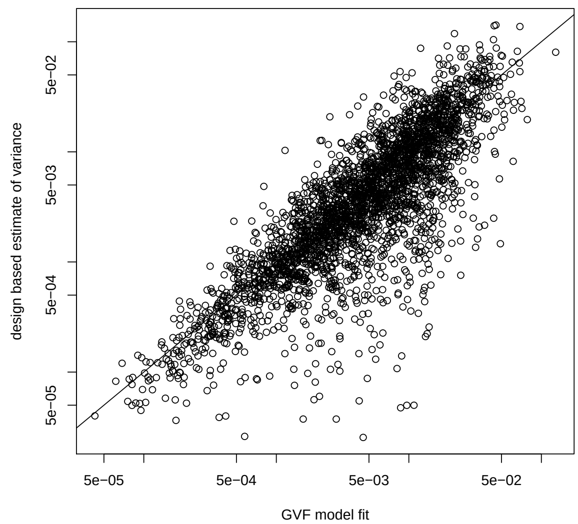

Figure 6: Ratio of GVF to Direct

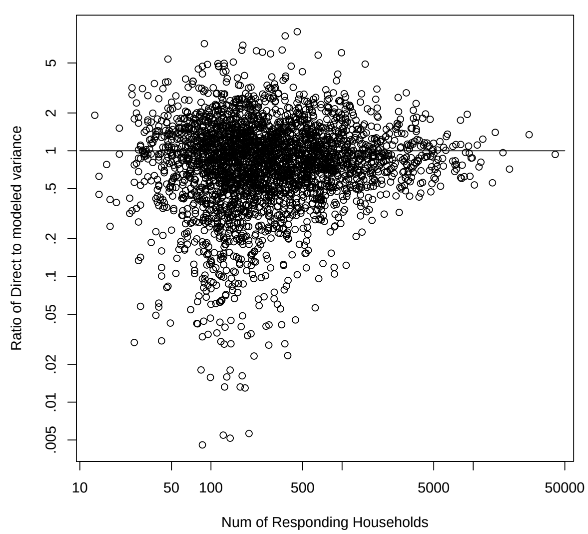

Figure 7: Ratio of GVF to Direct - Large Counties

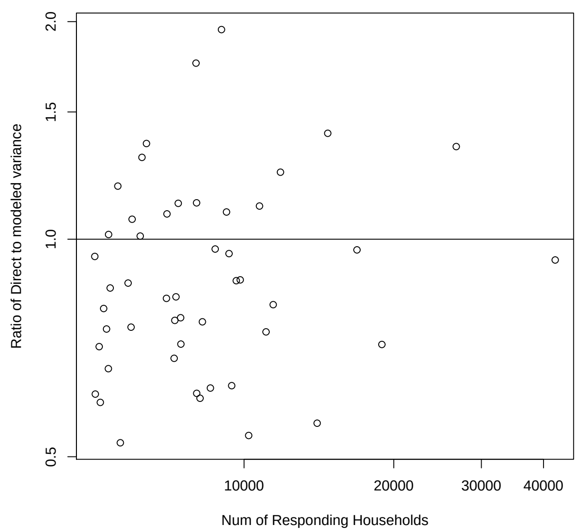

Figure 8: Estimated Degrees of Freedom

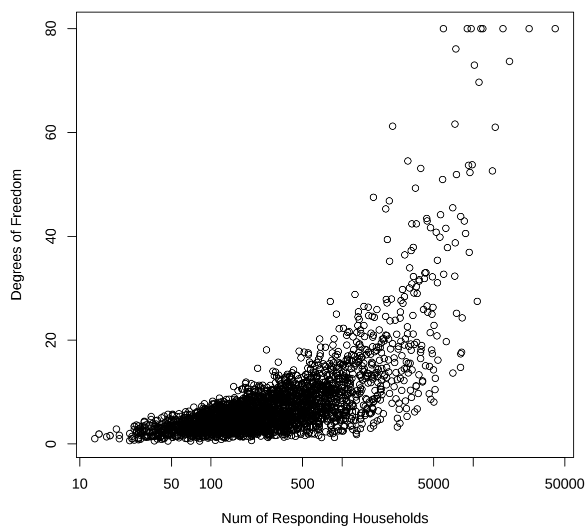

Figure 9: Smoothing Weights on Direct Variance Estimate

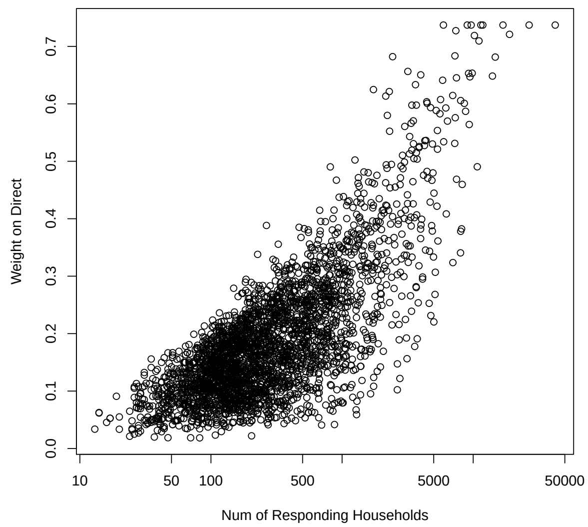

Figure 10: Distribution of Smoothing Weights

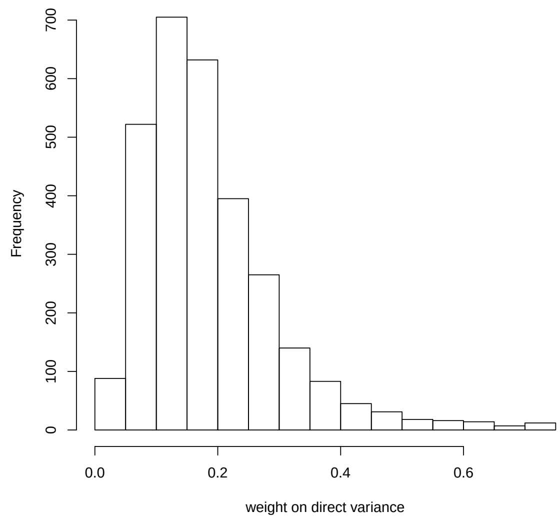

Figure 11: Empirical Bayes estimate of Variance

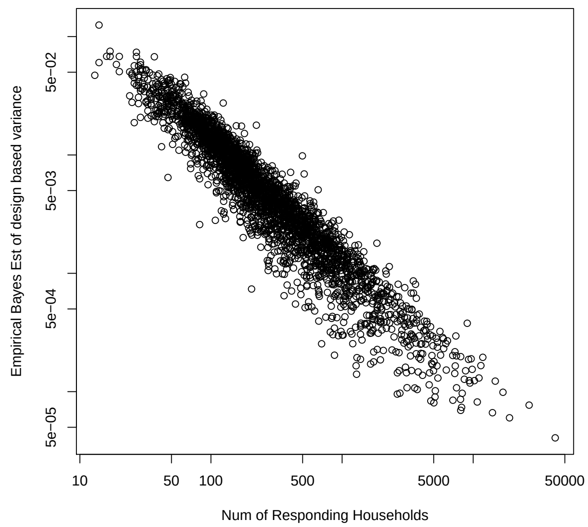

Figure 12: Ratio of Empirical Bayes to Direct

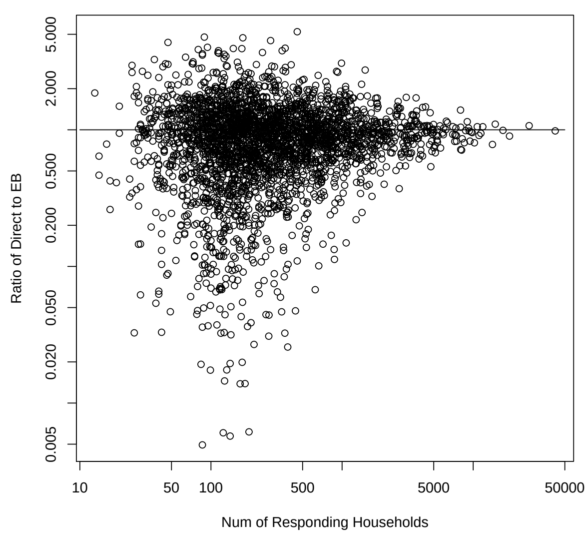

Figure 13: Ratio of Empirical Bayes to Direct - Large Counties