Imbedding Model-Assisted Estimation into ACS Estimation

Robert Fay
U.S. Census Bureau, 7252 Greentree Rd, Bethesda, MD, 20817

Abstract

The American Community Survey (ACS) began full implementation in 2005. Estimation for ACS, as typical for other large Census Bureau surveys, uses a complex series of ratio estimation and other adjustments to the weights. Working originally from ACS data for 1999-2001 in 36 test counties, previous research suggested that multiyear tract-level estimates could be improved by imbedding a step of model-assisted estimation, specifically generalized regression estimation (GREG), in the current ACS estimation. In particular, the GREG step incorporates administrative record data. Most data sets produced for the ACS Multiyear Estimates Study incorporate the GREG step, but alternative sets of 3- and 5-year estimates without the GREG step were produced for purposes of comparison. The paper will describe new refinements in the estimation approach and its potential future role in ACS estimation.

1. Introduction

In 2006, the American Community Survey (ACS) began its second year of full implementation. Based on the 2005 ACS data, the Census Bureau has published extensive tabulations for the household population for states, counties, places, and other geographic entities with population 65,000 or more (available at www.census.gov). Starting in the summer of 2007, a similar set of 1-year estimates will be published based on the 2006 data, this time incorporating data for major segments of the group quarters population. ACS plans call for continual annual publication of estimates based on single years of data for large geographic units above the population threshold of 65,000.

To reach smaller levels of geography, the ACS plans to publish 3-year period estimates based on combining data across 3 years, down to a threshold of 20,000 persons in the geographic unit. The first such national series will be available in 2008 for 2005-2007. But the ultimate objective is for the ACS to replace the decennial census long form, which in recent censuses has sampled households at approximately 1-in-6 overall. To provide the same level of geographic detail as previously afforded by the long form, the ACS plans to accumulate 5 years of data to create 5-year period estimates. The first such 5-year period estimates for the fully implemented ACS will become possible for the period 2005-2009, which is slated for publication in 2010.

This paper is one of a series of papers (Fay 2005a, 2005b, 2006) directed toward the development of statistical methods to improve the reliability of ACS sub-county estimates, particularly the 3-year and 5-year period estimates. Although the basic problem will again be summarized here, readers of the earlier papers will find its description virtually unchanged. The basic methodological approach, also summarized here, was sketched in Fay (2005a) and elaborated in Fay (2006) with the same basic features used in the current study. Unlike these earlier papers, however, this paper will set the ongoing research in the context provided by the recent report from the National Academy of Sciences (National Research Council, 2007) on the usability of the ACS. This paper also reports empirical findings based on new data. Specifically, the empirical analysis reported here is a component of the Multiyear Estimates Study, which produced 1-, 3-, and 5-year estimates based on ACS data in 34 test counties. Since 1999, the 34 counties had been sampled at rates either approximating or higher than the sampling rates for the national ACS implementation starting in 2005. Thus, the results in the 34 counties provide a preview of the future ACS.

The National Academy (National Research Council, 2007) identified major advantages of the ACS strategy over the decennial long form in terms of data quality, frequency, and timeliness. But the report also expressed concern over the lower reliability—that is, higher sampling variances—for the 5-year estimates relative to the decennial long-form estimates that they will replace. Consequently, the Academy's report emphasized the importance of research to reduce the variance of estimates and noted this research project on generalized regression estimation (GREG).

In short, the new results support the use of GREG estimation as a potential method to reduce variance of sub-county estimates, both for 3- and 5-year products. Key results will be presented in this paper, and Starsinic and Tersine (2007) present a more detailed analysis.

This paper targets two audiences. The first audience, of course, comprises my statistical peers in survey research. On the other hand, it is not intended to be accessible to the majority of ACS users, who generally would have limited interest or need to know these technical details. Instead, I hope I have made the key features of the new method clear to a second audience—the more technically sophisticated users of the ACS data, such as demographers, economists, and sociologists. I have observed that some members of this important group may have learned from past experience to suspect claims made by statisticians without the support of empirical data. I hope that the results presented here provide much of the evidence that these researchers might require, and that the paper forms a framework for additional analysis from the Multiyear Estimates Study.

On the other hand, the paper also shows why members of this important second audience of sophisticated users may freely ignore the presence of the GREG step in working with ACS data. In short, the GREG step does not change the basic frequentist paradigm on which the basic products from the ACS are based. Thus, users may interpret confidence intervals and standard errors reported for ACS estimates incorporating GREG estimation in the same way as if the GREG step were not present. Consequently, the presence of GREG estimation does not impose any new learning requirements on users who are comfortable working with the reported measures of reliability.

2. The Problem and Basic Approach

The Census Bureau divides counties into tracts, with an average population of roughly 4000 persons, and block groups within them, with an average population of roughly 1500 persons. In Census 2000 long-form estimates were published at both of these levels, with reduced detail at the block-group level. Tracts and block groups are among the lowest levels of geographic detail published in Census 2000 and previous censuses. The ACS will also publish estimates at the tract and block-group levels, as five-year period estimates, based on a designated sample of approximately 1-in-8 households built up from the monthly designated samples of approximately 1-in-480.

As previously noted, the ACS has maintained sampling rates in 34 test counties approximating or exceeding the sampling rates used for the full ACS in 2005. The first data products from this effort were 3-year period estimates for 1999-2001. Paul Voss and his colleagues (Van Auken et al. 2004, 2006) first noted that the tract-level sampling variances for ACS estimates for this period are considerably larger than initially projected, whereas county-level variances generally meet design predictions. Their analysis was based on the two Wisconsin test counties they studied, but Starsinic (2005) replicated this finding in the remaining 32 test counties. Differences between census and ACS estimation provided an explanation. In the census, raking-ratio estimation was performed within weighting areas that are typically coterminous with tracts, reducing to zero the variance of estimates controlled to known totals from the 100% census and benefiting other estimates, particularly of totals. In contrast, ACS weighting used controls at the county level and higher, because during the decade no controls are available at the tract level.

Fay (2005a) presented a preliminary argument that administrative record data could be combined with GREG estimation to increase reliability at the tract level for key estimates. The Census Bureau's program for administrative record data integrates data from a number of sources into a data file offering a census-like picture of the population. Specifically, the GREG estimation uses an extracted file of person-level records with assigned values for age, sex, race, and Hispanic origin, linked to addresses in the Census Bureau's Master Address File (MAF) when the available address information is sufficient. Not all administrative records can be linked to the MAF, however, and the coverage of the administrative records is less complete than the census. The completeness of coverage appears to vary geographically and, to some extent, by year. The GREG implementation to be described here is comparatively insensitive to these fluctuations, because the fluctuations can affect the variance improvements but are not a source of bias.

As previously outlined (Fay 2005a), the basic elements of the implementation are:

1. Link administrative records to the ACS sampling frame (the MAF), dropping administrative records that cannot be linked.
2. Form unweighted tract-level totals of the linked administrative record characteristics.
3. Apply ACS sampling weights at the housing-unit level to the linked administrative record data that fall into the ACS sample. The weighted estimates at this step represent unbiased (or essentially unbiased) estimates of the unweighted totals in step 2.
4. Using generalized regression estimation (GREG), calibrate the ACS sample weights so that the weighted administrative totals from the sample match the unweighted totals from step 2. (The number of constraints is allowed to vary with the size and other characteristics of each tract.)
5. Use the new housing-unit weights in subsequent stages of the ACS weighting, which includes ratio and raking/ratio estimation. Although the subsequent estimation steps adjust the new weights, the argument is that most of the variance reduction at the tract level will be retained in the final weights.

These steps will be described in more detail in subsequent sections.

3. Generalized Regression Estimation

The literature on generalized regression estimation (GREG) is extensive, and a relatively complete survey would require, at a minimum, a full-length review paper such as Fuller (2002). This paper will adopt the strategy of the previous paper (Fay 2006) in citing a few key references and pointing to more detailed reviews in Fay (2005a, 2005b).

Model-assisted estimation (Särndal, Swensson, and Wretman, 1992) includes both ratio estimation, familiar to almost all survey practitioners, and GREG as special cases. Many applications of GREG also qualify as a form of calibration estimation (Deville and Särndal 1992), and the class of calibration estimators offers some possible alternatives to the approach taken here. The discussion section will return to this point.

Previous papers employed a notation for the GREG attempting a compromise between the notation used by Särndal, Swensson, and Wretman (1992) and notation of Bankier and his colleagues (Bankier, Rathwell, and Majkowski, 1992; Bankier, Houle, and Luc, 1997; Bankier and Janes, 2003) in describing the use of GREG for estimation in the Canadian censuses of 1991, 1996, and 2001. The ACS application parallels several aspects of the Canadian application.

Instead, this paper will describe GREG in the notation used by Rao (2003, ch. 2) in his book *Small Area Estimation.* Although most of his book describes model-based indirect estimates, his chapter 2 focuses on direct estimators, including GREG. He remarks (2003, p. 10), "Effective use of auxiliary information through ratio and regression estimation is also useful in reducing the need for indirect small area estimators (Sections 2.3-2.5)." Indeed, two decades earlier Särndal (1984) similarly argued that model-assisted estimation might be a suitable choice for some small domain estimation problems. Fay (2005a) reviewed the theoretical literature in somewhat more detail as well as the evolving application to the Canadian census; the review in Fay (2005b) emphasized precedents among U.S. applications.

Consider a population U with values $y_1,\ldots,y_N$. Consider the estimation of the population total $Y = \sum_U y_j$ based on a sample s drawn with probability $p(s)$. Let $w_j(s) = w_j$ denote initial weights, either based on the inverse probability of selection, $w_j = \pi_j^{-1}$ or, more generally, weights based on π_j^{-1} adjusted by some early steps of estimation. Then $\hat{Y} = \sum_s w_j(s)\, y_j$ denotes the initial estimate of Y. (In the ACS application, the $w_j(s)$ include household noninterview adjustments, as described in Fay 2005a.) Suppose there are auxiliary data on p variables, with a known population total $X = (X_1,\ldots, X_p)^T$, where the superscript T denotes transpose. Suppose further that the corresponding individual vector x_j is known for each sample unit $j \in s$. (In the ACS application, the vector x_j is in fact known for every $j \in U$.)

An expression for the GREG estimator is given by

$$\hat{Y}_{GR} = \hat{Y} + (X - \hat{X})^T \hat{B} \qquad (1)$$

where $\hat{X} = \sum_s w_j\, x_j$ and

$$\hat{B} = (\hat{B}_1,\ldots,\hat{B}_P)^T = \left(\sum_s w_j\, x_j x_j^T / c_j\right)^{-1} \sum_s w_j\, x_j\, y_j / c_j \qquad (2)$$

for constants $c_j > 0$. (In the ACS application, $c_j = 1$, but other choices are possible.)

If $\hat{Y}_{GR}$ is expressed in terms of revised weights, $w_j^* = w_j^*(s) = w_j(s)\, g_j(s)$, then the "g-factor" or "g-weight"

$$g_j(s) = 1 + (X - \hat{X})^T \left(\sum_s w_j\, x_j\, x_j^T / c_j\right)^{-1} x_j / c_j \qquad (3)$$

is independent of the choice of Y. Consequently, the GREG step can be implemented as the weighting factor given by eq. *(3)*.

3.1 Ratio Estimation to the Frame Total

In the ACS case, a conceptually important special case arises for $p = 1$, $x_j = 1$. In effect, the constant term indicates membership in the ACS frame. Then for $c_j = 1$, eq. *(3)* reduces to

$$g_j(s) = 1 + (X - \hat{X})^T / \hat{X}^T = X / \hat{X}^T = X / \hat{X} \qquad (4)$$

where the transpose operation on scalars has no effect here with $p = 1$. In this case, GREG implements a simple ratio adjustment of the ACS sample weights to the ACS unweighted frame total, X. Previous results (Fay 2006) showed a substantial variance reduction in the estimated number of housing units when this estimator was applied at the tract level. The results to be reported here all incorporate 1 as the first element of any vector x,. But for almost all tracts $p > 1$, with the remaining elements of x_j based on administrative records.

Rao (2003, p. 13) notes that the weights given by eq. *(3)* calibrate the weighted estimate from the sample to the unweighted population total.

$$\hat{Y}_{GR}(x) = \sum_s w_j^*\, x_j = X \qquad (5)$$

GREG is a member of the class of calibration estimators (Deville and Särndal 1992). Fay (2006) reviewed GREG from the calibration perspective in more detail.

3.2 Asymptotic Unbiasedness of the GREG.

Unless there were very few observations in the calculation in eq. *(4)*, most survey researchers readily implement ratio estimation without hesitation in the appropriate context, and they worry little that its use would lead to appreciable bias. Of course, eq. *(4)* is not in general purely design-unbiased, because it is a non-linear expression. With an appropriate asymptotic framework, however, one can show that it is *asymptotically* unbiased, so that its mean square error is effectively given by its sampling variance. In the more general univariate case, ratio estimators of the form eq. *(3)* are asymptotically unbiased if $\hat{X}$ is unbiased for X.

Under some circumstances, there may be grounds to question whether $\hat{X}$ is unbiased for X in eq. *(4)*. When X is obtained from an independent source but the x_j are measured for the sample only, undercoverage of the survey frame or measurement error in x_j may affect the relationship. In many cases, the ratio estimate may still be desirable, but assessing its true impact is more complex than simply estimating its variance. In some cases, ratio estimation may reduce bias, but in other cases ratio estimation to an inappropriate X may increase it.

Analogously, the GREG estimator can be considered asymptotically unbiased if $\hat{X}$ is unbiased for X in eq. *(1)*. In some applications of GREG, this requirement may fail for the same reasons as in the univariate case (i.e. ratio estimation). As will be described in more detail, in the ACS application both $\hat{X}$ and X are based on identical information based on administrative record data matched to the ACS frame (the MAF). (A relatively small assumption is made about the ACS noninterview adjustments.) Thus, the ACS application assures that the requirement is met or essentially met through construction. Similarly, Statistics Canada has tried to minimize any differences between their sample and complete count data for characteristics used in their GREG estimation in the census. In other applications, however, there may be no basis to claim that $\hat{X}$ is unbiased for X in eq. *(1),* and issues of bias would then arise.

4. Methods

As previously noted, the ACS has been conducted in 34 test counties since 1999 at sampling rates approximating those in full production. In fact, rates during 1999-2001 were high enough so that the 3 years of data collected in the 34 counties were roughly comparable to the sample size that will be obtained from 5 years of the full ACS production. Until the Multiyear Estimates Study, the only 3- or 5-year period estimates that had been published from the ACS were those for 1999-2001.

The previously reported empirical research on GREG (Fay 2006) used the 1999-2001 data in conjunction with administrative records for year 2000 only. The results were quite encouraging, but several questions remained, including how well the procedure would work with administrative records updated annually, processed at a time removed from a census year.

This year, results from the Multiyear Estimates Study have been released (www.census.gov). Other Census Bureau researchers (Tersine and Asiala 2007; Starsinic and Tersine, 2007; Beaghen and Weidman, 2007) are documenting aspects of the study for the 2007 Joint Statistical Meetings. In addition to 1-year estimates previously available, 3-year period estimates are available for 1999-2001, 2000-2002, 2001-2003, 2002-2004, and 2003-2005. Similarly, 5-year period estimates are available for 1999-2003, 2000-2004, and 2001-2005. All published multiyear sets incorporate GREG estimation.

In conjunction with the Multiyear Estimates Study, 3-year period estimates for 2003-2005 and 5-year period estimates for 2001-2005 were prepared internally by excluding the GREG step. Because the GREG step is designed as an "imbedded" step that could either be implemented or skipped, it was comparatively easy to prepare versions without the GREG step for comparison.

Tersine and Asiala (2007) will describe the multiyear estimation in detail, but published descriptions for 1-year estimations (National Research Council, 2007; U.S. Census Bureau, 2006) provide accounts of the basic ideas. Estimation is a multistep process, with nine steps in the National Research Council's summary (2007, p. 5-17). The first four steps establish preliminary housing unit weights and adjust for household noninterviews. Although the noninterview adjustments imbed assumptions about nonresponse, the high weighted response rate to the ACS (approximately 97%) reduces the impact of the assumptions. The fifth step calculates and applies housing unit control factor 1, and in some counties the adjustment can be substantial. Consequently, the GREG estimation step was imbedded between the fourth and fifth steps. After the fifth step, ACS weights reflect an adjustment for, in effect, frame undercoverage. The adjustment for frame undercoverage invalidates the assumption that $\hat{X}$ is unbiased for X in eq. *(1)*.. In fact, if the GREG were applied after the fifth step, it would tend to undo the adjustment for frame undercoverage.

4.1 Choice of *X*

The same variables were used for both the 3-year and 5-year estimates. Whenever GREG was implemented, the first variable was identically 1 to always control to the frame total, that is, the ratio estimator described in the previous section. In a few small tracts, the algorithm actually controlled only to the frame total, but in almost all cases some combination of age/sex cells were implemented. In all cases, the age/sex data were based on the administrative record data summed to the housing-unit level and not any ACS data. Similar to Fay (2006), the age distribution was divided into the broad groups 0-17, 18-29, 30-44, 45-64, and 65+. Up to 3 alternative sets of age/sex cells were considered. The first set based on 7 cells plus the constant term was:

x1 = 1
x2 = 0-17 M+F,
x3 = 18-29 M+F
x4 = 30-44 M
x5 = 30-44 F
x6 = 45-64 M
x7 = 45-64 F
x8 = 65+ M+F;

a reduced 4-variable alternative was based on

x1 = 1
x2 = 0-17 M+F,
x3 = 18-44 M+F
x4 = 45+ M+F

and a 2-variable regression was based on

x1 = 1
x2 = total admin persons.

Note that all three equations calibrate to total administrative record persons. If the three regressions were each unacceptable, then the 1-variable regression used only

x1 = 1

to constrain the weights to agree with the number of units in the frame. This equation was the one previously discussed for its equivalence to ratio estimation.

The regression with the largest number of variables was used subject to the constraints that a true inverse could be computed in eq. *(2)* and that all $w_j^*(s) = w_j(s)g_j(s) > 0$.

Following the determination of age/sex variables in a given tract, the algorithm next determined if race/ethnicity variables could be included. The administrative record file contains designations for the race and ethnicity of the administrative record persons. In most cases, the values had been obtained through a match to Census 2000, but in others the values were derived from other sources, including imputation. For purposes of GREG, the information was combined into four groups: (1) Hispanic, (2) Black Non-Hispanic, (3) White Non-Hispanic, and (4) Other Non-Hispanic (including American Indian, native Alaskan, Asian, and Pacific Islander). An algorithm selected the largest group from (1), (2), or (4) such that at least 300 persons were in the selected group and that a residual group have at least 300 persons. The algorithm continued in this manner to identify potentially separate groups and then attempted to add race/ethnicity indicators to each in addition to the selection of age/sex variables. The same two constraints of non-singularity and positive weights were applied as at the tract level for 5-year periods.

A detail of the implementation should be noted here. The ACS frame changes each year (in fact, twice per year), and the gradual evolution of the frame must be taken into account. Step 2 refers to unweighted totals from the frame, but in fact the entries in the frame are time-weighted when determining *X*. For example, a unit in the frame for only one year would receive a time weight of 1/5 in a 5-year period estimate, but a unit in the frame all 5 years receives a time weight of 1. Equivalently, the 5-year *X* was the average of the annual versions of *X* for the years in the period.

For the 5-year period estimates, the GREG was not implemented in tracts with < 300 units in the frame. For 2001-2005, for example, GREG was not implemented in 45 tracts out of 2270 in the study.

GREG estimation for the 3-year period estimates was implemented at the place or place/MCD-part (minor civil division) level for weighting areas of approximately 10,000 persons or more, in place of tracts. Where necessary, GREG was also implemented for a county balance when the weighting areas defined by places left a residual area in the county.

As previously noted, the Multiyear Estimates Study computed period estimates for the 2001-2005 and 2003-2005 periods with and without GREG. The study published profiles for a variety of geographic areas. The profile includes 397 lines with estimated totals, including some duplicate lines. The following analysis is based only on estimated totals, omitting medians, ratios, proportions, and duplicate lines. For estimated totals of 0, the ACS publishes a positive variance, but for purposes of the comparison of estimates with GREG to without GREG, all estimated totals of 0 are excluded from the analysis. (In all cases, the estimates with and without GREG were either both zero or both positive.)

ACS variance estimation employs 80 replicate weights, which reflect all stages of estimation. Accordingly, the variance of GREG was reflected by implementing the GREG step for each of the 80 sets of replicate weights after the non-interview adjustments, and passing the replicate weights after GREG on to the next estimation step, the first step of housing unit ratio estimation. Because of the impact of subsequent estimation steps, the calibration effect of the GREG will be imperfectly preserved. But prior studies had already suggested that most of the variance gains from GREG at the subcounty level would be substantially retained, and this study provides a further empirical test of this claim.

5. Results

At the tract level, many estimates in the profiles are based on few ACS sample cases, so the corresponding sampling variances are large. To provide some perspective on tract-level reliability, Fig. 1 displays coefficients of variation (standard error as a percent of the estimate) by the size of the estimate. The figure shows some specific sets of estimates: (1) total housing units, which is substantially benefited from the control to the frame size; (2) total households; (3) total population; (4) 33 age/sex cells used by Starsinic and Tersine (2007), except for their use of median age; (5) 48 race/ethnicity cells, combining their 42 race cells and 6 Hispanic cells; and (5) 294 remaining cells of estimated totals in the profile, excluding duplicates.

Even in presenting interval averages, Fig. 1 shows the exceptionally wide variation in reliability depending on estimate size. Estimates below 300 appear quite unreliable for most purposes, those in the range 300-999 provide a rough indication of the true value, while those above 1000 frequently have cv's of 10% or less. Of course, estimates above 1000 generally represent a large proportion of the tract total, either on a person or a housing unit basis.

Suppose $\hat{Y}_{GREG}$ estimates a characteristic from the 2001-2005 period estimates including the GREG estimation step, and $\hat{Y}_{noGREG}$ the corresponding result when the GREG step is omitted. For a single tract, a natural comparison is the ratio of the final variances, $\mathrm{var}(\hat{Y}_{GREG})/\mathrm{var}(\hat{Y}_{noGREG})$, with values less than 1 reflecting positively on GREG. Some care is required in interpreting these ratios, however, because the variances are themselves estimates, and it is possible for the ratio under the null hypothesis of neutral GREG impact to have an expected value > 1. To summarize $\mathrm{var}(\hat{Y}_{GREG})/\mathrm{var}(\hat{Y}_{noGREG})$ over a set S of tracts, a set of variables, or a set of tracts crossed by variables, three statistics can be considered:

1. the ratio $\sum_S \mathrm{var}(\hat{Y}_{GREG})/\sum_S \mathrm{var}(\hat{Y}_{noGREG})$ of their totals.
2. their mean, $mean_S(\mathrm{var}(\hat{Y}_{GREG})/\mathrm{var}(\hat{Y}_{noGREG}))$; or
3. the median of the ratio of the relative variances, $med_S((\mathrm{var}(\hat{Y}_{GREG})/\hat{Y}_{GREG}^2)/(\mathrm{var}(\hat{Y}_{noGREG})/\hat{Y}_{noGREG}^2))$.

All three measures will be used in the comparisons. The average $(\hat{Y}_{GREG}+\hat{Y}_{noGREG})/2$ is used to classify estimates by size.

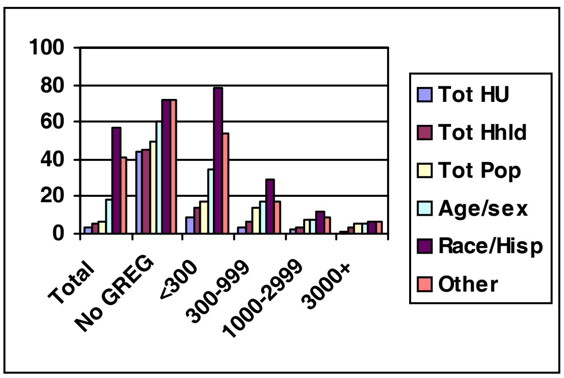

Fig. 1 Estimated percent coefficient of variation (cv) grouped by size of estimate. The results are at the tract level in 34 ACS test counties for the 5-year period 2001-2005 and include the GREG step. The average for all estimates is shown, as well as the cv's for the 45 smallest tracts in the study where GREG was not implemented. The larger cv's for race or ethnicity variables reflects the particularly high correlation of these variables within household.
Source: U.S. Census Bureau, Multiyear Estimates Study

Fig. 2 presents the variance comparisons for total housing units. For the 45 tracts where GREG was not implemented, no gains are expected, nor do they materialize. In all other cases, the variance reductions are dramatic.

Fig. 3 shows a similar comparison for total population. The results for total households are not shown, but they follow a pattern similar to Fig. 1 and 2. Unlike Fig. 2 and 3, Fig. 4 for age/sex cells and Fig. 5 for race and ethnicity cells both show that the choice of measure can affect the overall summary. But within given size intervals, such as 1000-2999 or 3000+, the three measures closely agree.

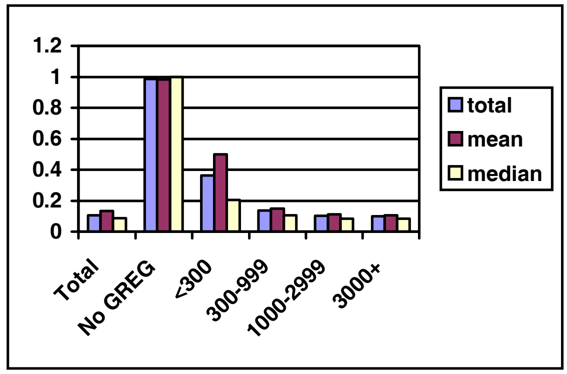

Fig. 2 Variance comparisons for total housing units at the tract level, 2001-2005. The measures shown are the ratio of the totals, the mean ratio, and the median rel-variance. For purposes of comparison, the cell < 300 is included, but it is almost empty in this case. Except for the 45 tracts where no GREG was attempted, the results all show substantial reductions in variance.
Source: U.S. Census Bureau, Multiyear Estimates Study

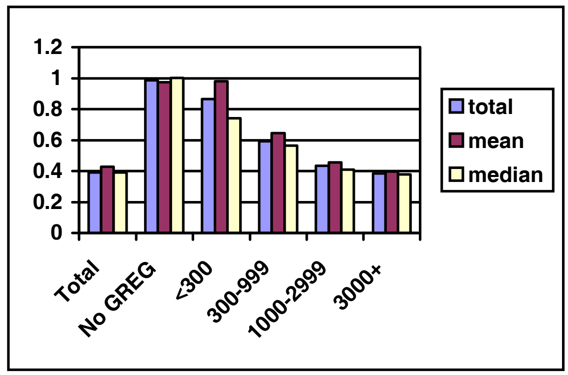

Fig. 3 Variance comparisons for total persons. As in Fig. 2., the cell < 300 is included, but it is again almost empty. The results consistently show approximately a 60% decrease in variance overall, although somewhat less in the smallest tracts.
Source: U.S. Census Bureau, Multiyear Estimates Study

In the long-form weighting for Census 2000 and previous censuses, raking-ratio estimation was implemented separately in a set of weighting areas, which were often tracts. Even though the weighting areas did not always precisely map onto other geographic entities such as places, they acted as building blocks to improve the precision of higher level aggregates. In the same way, implementation of GREG at the tract level affects the estimates of places and other sub-county geography. Fig. 7 parallels Fig. 6 in summarizing variance comparisons, this time at the sub-county place level. The relative gains are somewhat less than at the tract level, but they are nonetheless substantial.

As noted in the Methods section, a strategy of place/MCD parts was used to form weighting areas for the 3-year estimates, in place of tracts. Fig. 8 displays results at the place level for places meeting the 3-year threshold. For these larger places, the gains are somewhat less than in Fig. 7, nonetheless the ratio of total variances indicates a clear advantage to the GREG step.

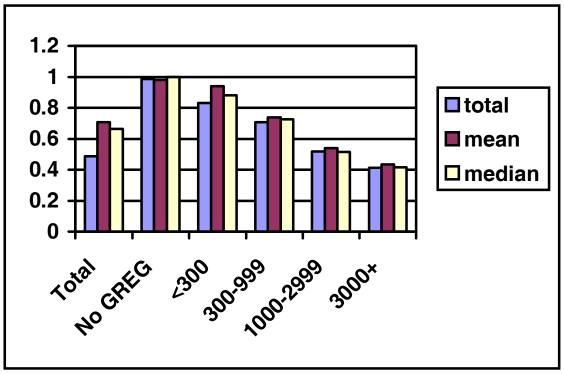

Fig. 4 Variance comparisons for 33 age/sex cells. More than in Fig. 2 and Fig. 3, the results summarized to a total level depend on the choice of measure. The overall average variance reduction is about 51% based on the ratio of totals, but the mean ratio indicates only 29% and the median 34%. Within ranges, however, the measures are more consistent, and all three would place the average variance reduction for estimates of 1000+ above 40%.
Source: U.S. Census Bureau, Multiyear Estimates Study

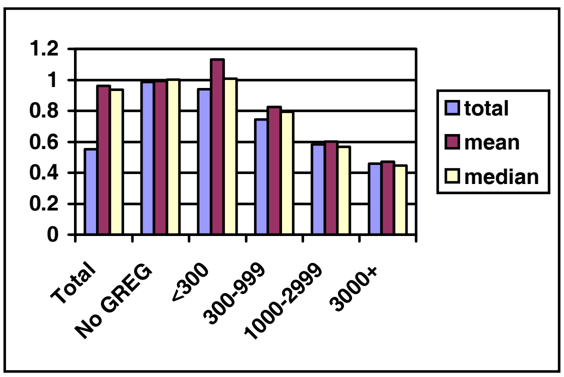

Fig. 5 Variance comparison for 48 race/Hispanic cells. As in Fig. 4, the comparison is sensitive to the choice of measures. The ratio of totals indicates an overall reduction of 45%, while the mean gives 4% and the median 6%. If restricted to estimates 1000+, all three measures put the average reduction around or above 40%.
Source: U.S. Census Bureau, Multiyear Estimates Study

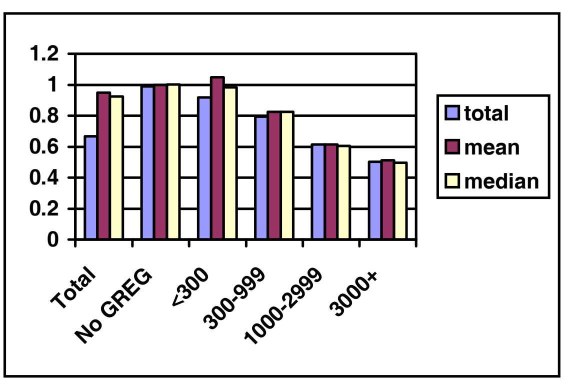

Fig. 6 Variance comparison at the tract level for the remaining 294 estimated totals in the profile. As in Fig. 3, 4, and 5, the comparison is sensitive to the choice of measures. The ratio of totals indicates an overall reduction of 33% compared to 5% for the mean and 7% for the median.
Source: U.S. Census Bureau, Multiyear Estimates Study

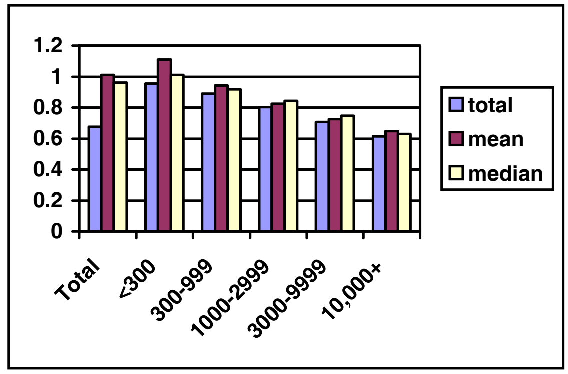

Fig. 7 Variance comparison at the place level for all 378 estimated totals in the profile, excluding duplicates, for the 5-year period 2001-2005. Places equivalent to complete counties: San Francisco, CA; Bronx, NY; and St. Petersburg, VA, are excluded.
Source: U.S. Census Bureau, Multiyear Estimates Study

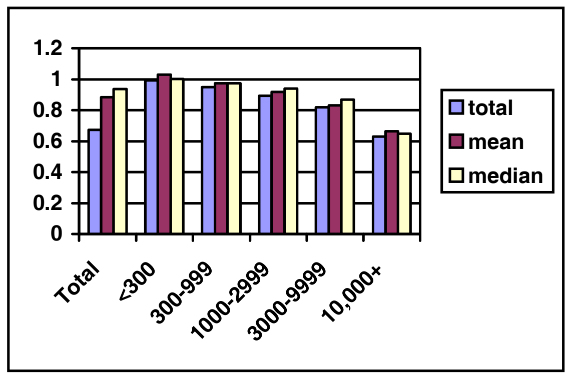

Fig. 8 Variance comparison at the place level for all 378 estimated totals in the profile, excluding duplicates, for the 3-year period 2003-2005. Only places meeting the ACS publication threshold of 20,000 population are included. Places equivalent to counties are omitted.
Source: U.S. Census Bureau, Multiyear Estimates Study

6. Discussion

As the title emphasized, the paper investigated the effect of *imbedding* GREG into the ACS estimation system. Although the GREG step precedes ACS ratio and raking steps that alter the impact of the GREG somewhat, the comparisons here based on the final weights (and replicate weights) show that the sub-county variance improvements from GREG survive largely intact.

A clear empirical finding is that the variance reductions from the GREG step differentially benefit larger estimates. Figure 1 shows that at the tract level the ACS profiles will contain a substantial number of noisy estimates, which the GREG at best only marginally improves. The largest relative gains from GREG occur for the larger estimates. In trying to summarize the impact of GREG, the three measures examined often differ. The measure based on the ratio of total variances places more weight on the variance reduction for larger estimates, because the larger estimates tend to have larger variances, particularly without GREG. Global measures based on the median or mean ratio place equal emphasis on all tracts.

ASA page restrictions effectively limit the number of analyses that can be included in this paper, and clearly more detailed comparisons from these data are of interest. The results here complement other analyses, including Starsinic and Tersine (2007).

Some specific follow-on analyses may be of particular interest to the sociologists and other social scientists considered to be a potential audience for this paper. Some of them may be concerned that the introduction of the GREG into ACS might systematically shift the estimates in some manner, indicative of a bias in spite of the theory presented in Section 3. Because of ratio controls, it can be noted that whether GREG is included or not doesn't change the estimated population or number of housing units at the county level or higher. A detailed presentation of the empirical data can be used to show that, although the remaining estimated totals change somewhat at the county level and higher, the net change is generally very small.

A parallel analysis at the tract level could be used to show that, although the tract level estimates do change somewhat at this level as a result of the introduction of GREG, the shifts are within reason considering the goal of variance reduction.

Further analyses could be used to suggest improvements to the GREG. For example, the Canadian census implements a two-tiered implementation of GREG (Fay 2005a).

For some large places, the 3-year estimates may have more ACS data available than tracts do in 5 years. A more sophisticate form of GREG, using additional variables or interacting the current ones, could yield further improvements to the 3-year estimates.

Finally, future research could cast a somewhat wider net of possible alternative estimators, possibly other forms of calibration estimation. A key feature of the GREG step in this application is that it adjusts a housing unit weight, which in turn is then adjusted by a housing unit ratio factor. Some forms of calibration estimation, such as the usual implementation of ratio-raking estimation, are not suited to the constraint of producing a single housing unit weight. Nonetheless, many of the alternative calibration estimators surveyed by Deville and Särndal (1992) would be appropriate for a single housing unit weight, and they could be considered.

Note: (1) This report is released to inform interested parties of ongoing research and to encourage discussion of work in progress. Any views expressed on statistical issues are those of the author and not necessarily those of the U.S. Census Bureau. I wish to thank Anthony Tersine and Michael Beaghen for comments on an earlier version of this paper.

References

Bankier, M.D., Rathwell, S., and Majkowski, M. (1992), "Two Step Generalized Least Squares Estimation in the 1991 Canadian Census," *Proceedings of the Survey Research Methods Section,* American Statistical Association, pp. 764-769.

Bankier, M., Houle, A.-M., and Luc, M. (1997), "Calibration Estimation in the 1991 and 1996 Canadian Censuses," *Proceedings of the Survey Research Methods Section,* American Statistical Association, pp. 66-75.

Bankier, M. and Janes, D. (2003), "Regression Estimation of the 2001 Canadian Census," *Proceedings of the 2003 Joint Statistical Meetings on CD-ROM,* American Statistical Association, pp. 442-449.

Beaghen, M. and Weidman, L. (2007), "Statistical Issues and Interpretation of the American Community Surveys One-, Three-, and Five-Year Period Estimates," prepared for presentation at the Joint Statistical Meetings, Salt Lake City, UT, 29 Jul-2 Aug 2007.

Deville, J. and Särndal, C.-E. (1992), "Calibration Estimators in Survey Sampling," *Journal of the American Statistical Association, 87,* 376-382.

Fay, R.E. (2005a), "Model-Assisted Estimation for the American Community Survey," *Proceedings of the 2005 Joint Statistical Meetings on CD-ROM,* American Statistical Association, Alexandria, VA, pp. 3016-3023.

______ (2005b), "Potential Applications of Model-Assisted Estimation to Demographic Surveys in the U.S.," paper presented at the Federal Committee on Statistical Methodology Research Conference, available from www.fcsm.gov/05papers/Fay_IIIC.pdf.

______ (2006), "Using Administrative Records with Model-Assisted Estimation for the American Community Survey," *Proceedings of the 2006 Joint Statistical Meetings on CD-ROM,* American Statistical Association, Alexandria, VA, pp. 2995-3001.

Fuller, W.A. (2002), "Regression Estimation for Survey Samples," *Survey Methodology, 28,* 5-23.

National Research Council (2007), *Using the American Community Survey: Benefits and Challenges, Panel on the Functionality and Usability of Data from the American Community Survey,* Citro, C.F. and Kalton, G. (eds.), Committee on National Statistics, Division of Behavioral and Social Sciences and Education, The National Academies Press, Washington, DC.

Rao, J.N.K. (2003), *Small Area Estimation,* John Wiley, New York.

Särndal, C.-E. (1984), "Design-Consistent Versus Model-Dependent Estimation for Small Domains," *Journal of the American Statistical Association, 79,* 624-631.

Särndal, C.-E., Swensson, B., and Wretman, J. (1992), *Model Assisted Survey Sampling,* Springer-Verlag, New York, NY.

Starsinic, M. (2005), "American Community Survey: Improving Reliability for Small Area Estimates," *Proceedings of the 2005 Joint Statistical Meetings on CD-ROM,* American Statistical Association, pp. 3592-3599.

Starsinic, M. and Tersine Jr., A. (2007), "Analysis of Variance Estimates from American Community Survey Estimates," prepared for presentation at the Joint Statistical Meetings, Salt Lake City, UT, 29 Jul-2 Aug 2007.

Tersine Jr., A. and Asiala, M. (2007), "Methodology for the Production of American Community Survey Multiyear Estimates," prepared for presentation at the Joint Statistical Meetings, Salt Lake City, UT, 29 Jul-2 Aug 2007.

U.S. Census Bureau (2006), "Design and Methodology—American Community Survey," Washington, DC: U.S. Department of Commerce, issued May 2006, from www.census.gov/acs/www.

Van Auken, P.M., Hammer, R.B., Voss, P.R., and Veroff, D.L. (2004), "American Community Survey and Census Comparison, Final Analytical Report, Vilas and Oneida Counties, Wisconsin; Flathead and Lake Counties, Montana," unpublished report dated March 5, 2004, available at http://www.census.gov/acs/www/AdvMeth/acs_census/lreports/vossetal.pdf.

______ (2006), "The American Community Survey in counties with 'seasonal' populations," *Population Research and Policy Review, 25,* 275-292.