

Two Simple Methods to Improve Official Statistics for Small Subpopulations

Nikolas Mittag*

CERGE-EI

September 5, 2018

Abstract: Many important statistics are known from official records for the entire population, but have to be estimated for subpopulations. I describe two methods that are straightforward to implement and can reduce the substantial sampling error of the commonly used direct survey estimates for small subpopulations. Both estimators are simple variants of (known) ways to increase accuracy by modeling time-series and cross-sectional variation in repeated surveys. The first estimator incorporates information from repeated cross-sections. The second estimator additionally uses the knowledge of the statistic for the overall population to improve accuracy of the estimates for subpopulations. To illustrate the performance of the estimators and their advantages for practitioners, I compare the estimated number of female and elderly recipients of a government transfer program by county to the “true” number from administrative data on all recipients in New York. I find that even the simple estimators substantially improve survey error. Incorporating the statistic of interest for the overall population yields particularly large error reductions and can reduce non-sampling error.

*Any opinions and conclusions expressed here are those of the author and do not necessarily represent the views of the New York Office of Temporary and Disability Assistance (OTDA) or the U.S. Census Bureau. The analyses in this report were conducted as part of a project of the U.S. Census Bureau Center for Administrative Records Research and Applications. The data was accessed at the Chicago Census Research Data Center and the results were reviewed to prevent the disclosure of confidential information. The data are confidential and cannot be published, but researchers can apply for access by writing a proposal to the U.S. Census Bureau. This research was supported by the Czech Science Foundation (16-07603Y), Charles University (UNCE/HUM/035) and the Czech Academy of Sciences (RVO 67985998). This report would not have been completed without the help of Trent Alexander, Pablo Celhay, Dave Dlugolecki, Graton Gathright, Bruce Meyer and Joey Morales. Patrick Gaulé, Štěpán Jurajda and two anonymous referees provided great comments. All errors are mine. Address: CERGE-EI, joint workplace of Charles University and the Economics Institute of the Czech Academy of Sciences, Politických vězňů 7, Prague 1, 111 21, Czech Republic, nikolasmittag@posteo.de

Keywords: survey error, small populations food stamps, government transfers, official statistics.

1 Introduction

Statistics such as unemployment and poverty rates or the number of recipients of government benefits for counties, cities or other small geographic areas are important indicators of local economic conditions and well-being. In the U.S., the Census Bureau and the Bureau of Economic Analysis (BEA) produce a wide range of such statistics. For example, the BEA provides statistics on income and transfer programs by county and the Census Bureau produces small-area estimates of poverty (SAIPE) and health insurance coverage (SAHIE). The same statistics are often of interest for subpopulations such as by race, gender or age as well. Among others, they are used to evaluate policies targeted at these subpopulations, to assess their economic conditions and to coordinate outreach efforts. For example, poverty rates are usually also estimated for groups with particularly high poverty risk, such as single parents or individuals with disabilities (e.g. U.S. Census Bureau, 2015). Participation in government transfer programs is often calculated for demographic subgroups to evaluate their access to government benefits (e.g. Haider, Jacknowitz and Schoeni, 2003).

Government agencies often publish reliable aggregate statistics for the entire population based on administrative records, but researchers and policy makers interested in subpopulations usually estimate these statistics from survey data. The main source of such estimates in the US is the American Community Survey (ACS). Yet even with the large sample size of the ACS, estimates for small areas such as counties are often imprecise, particularly for subpopulations. The literature on small domain estimation proposed methods to improve precision (see e.g. Rao and Molina, 2015, Pfeffermann, 2002, 2013). Nevertheless, many studies and government agencies rely on standard, direct survey estimates instead (e.g. Bohn et al., 2013, D’Onofrio et al., 2015, Cerf Harris, 2014). This reliance on standard methods is neither due to a lack of advanced small domain methods nor their performance. Programs like SAIPE and SAHIE underline that these methods work well in practice, but also that they require more complex models, methods or additional assumptions. Researchers likely often see standard survey estimates as less “costly”, possibly in terms of introducing flawed modeling

choices that could introduce bias, the cost of implementing them or their audience not being familiar with more complex methods. Practitioners, especially those who do not produce small domain estimates on a regular basis, could still benefit from parts of the substantial improvements of small domain methods by using simplified small domain methods. Many practitioners relying on direct survey estimates are experts in their substantive fields, rather than small domain estimation, so methods that make it straightforward to assess whether they are useful in a given case and are simple to implement could improve this situation.

I describe and evaluate two simple estimators to improve estimate accuracy for small subpopulations. The first estimator uses area- and time-fixed effects and thereby combines information from multiple areas and time periods to improve accuracy. The second estimator combines the survey data with aggregate statistics on the overall population to improve estimates for subpopulations. This population adjusted estimator is simple to implement and can also help to address other survey errors, such as underreporting. Both estimators can be seen as simple variants of methods from the previous literature. Their key advantages are that the estimators are simple to implement and that the conditions under which they improve precision are straightforward to assess empirically. Their performance and ease of implementation make them attractive for practitioners, but the methods discussed here should not be seen as substitutes for those willing to implement more complex estimators.

To illustrate these methods and evaluate how they can improve over the direct survey estimates, I estimate the number of recipients of the Supplemental Nutrition Assistance Program (SNAP) per county using the ACS public use data for two subpopulations, females and the elderly.¹ To assess the accuracy of the estimates, I compare them to the true number of recipients in these subpopulations according to administrative records. Data that provide such a measure of truth would solve the problem in practice, but are rarely available to researchers. In summary, the results suggest that even simple methods can substantially

¹SNAP is the largest means-tested government transfer program in the US, see Hoynes and Schanzenbach (2016) for a detailed description. SNAP receipt rates for females and the elderly are not available from official statistics and direct survey estimates are noisy.

improve the accuracy of official statistics for small subpopulations. The population adjusted estimator yields large error reductions in terms of mean squared error (MSE) or mean absolute deviations (MAD), cutting error by up to 77 percent compared to direct survey estimates. The findings also show that this estimator can reduce bias from non-sampling error such as underreporting. The error reductions of the fixed effects estimator are smaller, but it still improves over the survey estimates. Whether the methods yield error reductions of this magnitude more generally remains an open question. I line out conditions under which they are likely to do so. These conditions provide guidance to practitioners in their choice whether the methods are likely to improve precision in their application and which methods are particularly promising.

The remainder of this report proceeds as follows: Section 2 introduces the estimators. Section 3 describes the data. Section 4 presents results. Section 5 concludes.

2 Two Simple Estimators

Let R_{it} be an indicator that individual i receives SNAP at any point in year t , and let R_{it}^S be an indicator that individual i receives SNAP in year t and is in the subpopulation of interest. Let x_{ct} be the total number of recipients in county c in year t and \hat{x}_{ct} the direct survey estimate of this number. Similarly, let y_{ct} be the statistic of interest, the number of recipients belonging to the relevant subpopulation in county c in year t . \hat{y}_{ct} is the corresponding direct survey estimate. N_{ct} is the sample size for county c at time t , N_c and N_t are the sample sizes corresponding to county c and time period t respectively.

The two estimators I use both borrow strength over time and space, which makes them simple versions of more advanced methods. Pfeffermann and Burck (1990) propose a general time series model. Rao and Yu (1994) additionally consider borrowing cross-sectional information. Datta et al. (1999) extend this model and discuss the statistical properties of such estimators in a hierarchical Bayes framework. See Pfeffermann (2002) for an overview

of early theoretical contributions. For applications that extend this methodology see e.g. Pfeiffermann and Bleuer (1993), You, Rao and Gambino (2003) and You (2008). Statistical agencies such as the U.S. Census Bureau (Bauder, Luery and Szelepka, 2018, Luery, 2010), the U.S. Bureau of Labor Statistics (Tiller, 1992, Pfeiffermann and Tiller, 2006) and Statistics Netherlands (Bollineni-Balabay, van den Brakel and Palm, 2016, van den Brakel and Krieg, 2016) employ more advanced models to produce small area statistics on a routine basis.

The first estimator combines information from repeated cross sections. In particular, I use pooled cross-sections to estimate a simple two-way fixed effects model, i.e. an OLS regression² of benefit receipt on dummies for each county and time period:

$$R_{it}^S = \theta_c + \gamma_t + \varepsilon_{it} \quad (1)$$

I then estimate the number of recipients in county c at time t as the sum of the corresponding time and county fixed effect multiplied by the population size in county c in year t , P_{ct} :

$$\hat{y}_{ct}^{FE} = (\hat{\theta}_c + \hat{\gamma}_t)P_{ct} \quad (2)$$

This fixed effects estimator is an adjusted version of the common practice to pool multiple survey years (U.S. Census Bureau, 2008), which amounts to estimating area fixed effects. Adding time fixed effects additionally allows for annual shocks that are common to all areas. The efficiency gains of the fixed effects estimator stem from incorporating time-series information from other editions of the sample as well as information from other areas. Rather than using N_{ct} observations to estimate $\mathbb{E}[y_{ct}]$ as $\hat{\theta}_{ct} + \hat{\gamma}_{ct}$, one uses N_c observations to estimate $\hat{\theta}_c$ and N_t observations to estimate $\hat{\gamma}_t$. This implicit increase in sample size yields smaller gains as the sample gets larger. Thus, one would expect this estimator to work particularly well for particularly small populations. Compared to simply pooling years, the

²One may often prefer to use a non-linear link function. The model does not include continuous covariates and the predicted probabilities in the application are all far from 0 or 1. Therefore, results using standard link functions are similar to the OLS results.

estimator will yield particularly large variance reductions when the variance of $\hat{\gamma}_{ct}$ is large relative to the variance of $\hat{\gamma}_t$, which will be the case whenever (common) time-specific effects matter, but are imprecisely estimated at the small area level.

Pooling years assumes that there are no time effects, but the estimator above allows for shocks to the participation rate that are common to all counties. Nevertheless, it still assumes that there are no county specific time effects. This assumption should be tested in the two-way ANOVA. It is violated, for example, if there are county specific time trends in SNAP receipt, which would cause bias. Yet, the estimator still improves MSE relative to the direct survey estimator as long as the squared bias is smaller than the reduction in variance, i.e. when the variance of the county-specific shocks is small relative to the variance of the sampling error, which seems likely for short time periods such as the one considered here. If the ANOVA indicates that time-effects may differ between counties or to further increase accuracy one can model either the cross-sectional or time-series variation better. A simple way to do so is to incorporate covariates or allow for less aggregated time effects. See e.g. Datta et al. (1999) for more advanced methods to do so and Bollineni-Balabay, van den Brakel and Palm (2016) for a specific application. Datta et al. (1999) also derive the statistical properties of such more advanced methods. The fixed effects estimator is a standard OLS predictor, so its bias, variance and MSE follow from standard results.

The second method combines survey data with information on the same statistic for the entire population. It makes use of the fact that government agencies often publish the total number of recipients in the entire population, x_{ct} , but usually do not provide these statistics for subpopulations, so the researcher still has to estimate y_{ct} from survey data. The second estimator uses the survey error for the entire population $x_{ct} - \hat{x}_{ct}$ to improve the precision of the estimate for the subpopulation of interest. If sampling is not related to subgroup membership, the fraction of subpopulation members missed by the survey (i.e. “in $x_{ct} - \hat{x}_{ct}$ ”) is the same as the fraction of subpopulation members among all recipients

(in county c at time t), denoted by f_{ct} .³ Thus, the expectation of the sampling error in the survey estimate for the subpopulation, $y_{ct} - \hat{y}_{ct}$, conditional on the sampling error in the overall population, $x_{ct} - \hat{x}_{ct}$, is $\mathbb{E}(y_{ct} - \hat{y}_{ct} | x_{ct} - \hat{x}_{ct}) = f_{ct}[x_{ct} - \hat{x}_{ct}]$. Therefore, the following population adjusted estimator may yield efficiency gains:

$$\hat{y}_{ct}^{PA} = \hat{y}_{ct} + \hat{f}_{ct}[x_{ct} - \hat{x}_{ct}] \quad (3)$$

where \hat{f}_{ct} is an estimate of the fraction of recipients in the subpopulation. The estimator is consistent as long as \hat{f}_{ct} is consistent and becomes more efficient as the estimate of f_{ct} becomes more precise. Thus, one may want to borrow strength across space and time to use the largest sample for which f_{ct} is constant to estimate it. Since the demographic characteristics of SNAP recipients vary substantially between counties, but are relatively stable over short time periods, I estimate f_{ct} by pooling all years for county c . Thus, I allow f_{ct} to vary between counties, but assume that it does not change over time. One can use the more sophisticated models of time-series and cross-sectional variation cited above to relax this assumption or increase accuracy further.

In addition to borrowing strength across time and space, the population adjusted estimator uses information from administrative data in a simple way. It is a form of a generalized regression or calibration estimator (Deville and Särndal, 1992, Kott, 2009), similar to the ratio estimators in, for example, Harter, Macaluso and Wolter (2005), Pfeffermann (2002) and West (1983). It becomes a simple ratio estimator, $\hat{y}_{ct}x_{ct}/\hat{x}_{ct}$, if one estimates f_{ct} by $\hat{y}_{ct}/\hat{x}_{ct}$, i.e. if one does not exploit information from other areas or time periods to estimate f_{ct} . Zanutto and Zaslavsky (2002) also use administrative data to improve small area estimates, but do not use it as a constraint as the population adjusted estimator does. The literature proposed several more advanced methods to improve small domain estimates by combining data sources (e.g. Gee and Fisher, 2004, Kim, Park and Kim, 2015, Tarozzi and

³That is, f_{ct} is the probability that a recipient from county c is a subpopulation member, i.e. $\Pr(R_{it}^S = 1 | R_{it} = 1)$ for individuals in county c .

Deaton, 2009) or imposing aggregate constraints (e.g. Pfeiffermann and Tiller, 2006). Little (2012) specifically mentions incorporating administrative aggregates to improve small area estimates in a calibrated Bayes framework. One can also view the population adjusted estimator as a simple version of the method of moment estimator proposed by Imbens and Lancaster (1994) that imposes true total receipt as a constraint. Using results from this literature on calibration or method of moment estimators may also help to clarify the statistical properties of the estimator in finite samples, which are beyond the scope of this paper. The calibrated Bayes framework Little (2012) discusses provides a general framework to extend the methods and analyze their properties. If reliable measures of other related statistics, such as poverty or unemployment rates, are available, using them in a generalized method of moments estimator as in Imbens and Lancaster (1994) can be used to further increase accuracy. Harter, Macaluso and Wolter (1999) examine several other possible extensions.

The two sources of additional information also point to circumstances in which the population adjusted estimator should be attractive more generally. First, by using accurate information on the total number of recipients, the population adjusted estimator estimates the fraction instead of the number of recipients that are members of the subpopulation. Thus, one may expect improved accuracy when the estimate of the fraction is (relatively) more precise, such as when the fraction of subpopulation members is close to 0.5. Second, the population adjusted estimator uses a more precise estimate of f_{ct} . Thus, one would expect the population adjusted estimator to work particularly well when borrowing strength across time or space increases the precision of f_{ct} as discussed above.

Taken together, one may expect the population adjusted estimator to work well more generally for larger subpopulations. Another common case in which one would expect the population adjusted estimator to work well is when the size of the overall population varies substantially across time or space, but the composition in terms of the subpopulations of interest is more stable. On the other hand, without any restrictions on the covariance of \hat{y}_{ct} , \hat{x}_{ct} and \hat{f}_{ct} , the MSE of the population adjusted estimator can be higher than the MSE

of the direct survey estimator even when both estimators are consistent. The population adjusted estimator can also increase MSE when its assumptions fail, i.e. if the estimate of f_{ct} is biased or the administrative information on the overall population size, x_{ct} , is wrong.

In addition to reducing sampling variation, the population adjusted estimator can address other sources of survey error such as underreporting. \hat{y}_{ct}^{PA} is consistent as long as the fraction of recipients in the subpopulation from the survey is a consistent estimator of the true fraction, i.e. as long as survey error does not differ systematically between the subpopulation and the population overall. On the other hand, if the coverage rates of the survey for the population ($c^x = \mathbb{E}(\hat{x})/x$) and the subpopulation ($c^y = \mathbb{E}(\hat{y})/y$) differ, \hat{f}_{ct} is inconsistent and the estimator is biased. Such differences in coverage rates could arise if one of the groups reports more accurately. Even if it is inconsistent, the estimator still reduces bias if the coverage rate of the subpopulation is not too much better than that of the overall population. In particular, if the survey underestimates receipt in the overall population ($x_{ct} > \mathbb{E}(\hat{x}_{ct})$), the estimator reduces bias as long as $\tilde{f}_{ct}[x_{ct} - \tilde{x}_{ct}] < 2[y_{ct} - \tilde{y}_{ct}]$.⁴ Re-arranging and using the definitions of f_{ct} and the coverage rates, this condition becomes

$$c^y < 2 \frac{c^x}{c^x + 1} \quad (4)$$

when $x_{ct} > \mathbb{E}(\hat{x}_{ct})$ and, analogously, when $x_{ct} < \mathbb{E}(\hat{x}_{ct})$

$$c^y > 2 \frac{c^x}{c^x + 1}.$$

These formulas show that if c^y is too close to one⁵ relative to c^x , i.e. when the coverage of the subpopulation is too good relative to the coverage of the overall population, this condition fails and the estimator overcorrects. Intuitively, the bias reductions will be particularly large

⁴Where the tilde in \tilde{f}_{ct} , \tilde{x}_{ct} and \tilde{y}_{ct} denotes the probability limit of the corresponding estimate (to allow for cases where the limit differs from the true value).

⁵In the extreme case where c^y is on the other side of one, because the survey overestimates one population and underestimates the other one, the condition never holds and the correction is in the wrong direction.

if the estimates for the overall population and the subpopulation over- or underestimate the statistic of interest by the same multiplicative factor, i.e. coverage rates for the overall population and the subpopulation are similar ($c^x \approx c^y$). Since c^x is simple to estimate in practice, one can assess whether the estimator reduces bias if information on the differences in coverage rates is available.

The population adjusted estimator is a non-linear combination of estimated quantities, so estimating variance and MSE using the delta method is complex. Re-writing the estimator as a method of moments estimator as in Imbens and Lancaster (1994) makes it simple to obtain standard errors even for stratified or non-iid samples. The moment conditions for each county and year are

$$\begin{aligned}\mathbb{E}(\tilde{y}_{ct}^{PA} - \tilde{y}_{ct} - f_c[x_{ct} - \tilde{x}_{ct}]) &= 0 \\ \mathbb{E}(\tilde{y}_{ct} - N_{ct}R_{it}^S) &= 0 \\ \mathbb{E}(\tilde{x}_{ct} - N_{ct}R_{it}) &= 0 \\ \mathbb{E}(f_c - R_{it}^S | R_{it} = 1) &= 0\end{aligned}\tag{5}$$

where the first three moments are for the population in county c at time t , while the fourth moment is for the population from which one estimates f_{ct} , i.e. all years for county c in the application here.⁶ Formulas for the standard errors then follow from the method of moments formulas as described in Imbens and Lancaster (1994) and are part of the standard output of method of moment routines of many software packages.

3 Data

I use three data sources to evaluate the estimators: administrative micro data on SNAP recipients, ACS survey data and aggregate data on SNAP receipt.

⁶For other choices, the subscript of f should be modified.

The administrative micro data are records of all SNAP payments in NY from 2007 through 2012 provided by the New York State Office of Temporary and Disability Assistance (OTDA). The records are from actual payments and contain identifiers for every individual on a SNAP case that have been checked by OTDA against social security records. The data appear to be accurate. For example, total payments differ from official aggregate outlays by less than a percent in all years. These data have been used in several papers that discuss their accuracy further (e.g. Meyer and Mittag, forthcoming, Celhay, Meyer and Mittag, 2018*a,b*, Cerf Harris, 2014, Scherpf, Newman and Prell, 2014). The records also contain simple demographic information that is rarely available in administrative data. In particular, having information on age, gender and addresses allows me to calculate the number of female and elderly (defined as those 60 and older) recipients by county.⁷

The data contain the universe of recipients and appear very accurate, so they provide me with a measure of “truth” for the number of recipients in each county. Thus, researchers could calculate the statistic of interest from these administrative records directly, but such records are usually not available to the public and are not required to implement the methods discussed here. Access to these administrative records enables me to evaluate the estimators described above by comparing the estimates to this standard of “truth”. In practice, one would use only survey data (and aggregate statistics) to implement the estimators I discuss.

The survey sample I use is the 2008-2012 ACS sample of individuals in NY. The ACS is the largest US household survey, including approximately 2.5 percent of the population each year. See U.S. Census Bureau (2014) for detailed information on survey design. The ACS is the most important source of statistics for small geographies, because of its large sample size. It collects detailed information on housing, income and program receipt. This information is commonly used to estimate statistics similar to the SNAP receipt rates in the

⁷The data are likely to contain a small number of errors due to misreported addresses or mobility. These errors should be negligible relative to the sampling variation of all estimators, particularly for the information on the county of residence, which is important for administrative purposes, as it needs to match the county of the responsible SNAP office. Comparing the county in the address with the information on the responsible SNAP office confirms the accuracy of the data.

application below. Thus, similar problems and solutions are likely to be of relevance more generally. Specifically, the sample for subpopulations in small areas is often still too small to yield reliable estimates. The Census Bureau recommends pooling multiple years of data when estimating statistics for areas with less than 65,000 people (U.S. Census Bureau, 2008). Thus, for about 75 percent of counties, even the largest US survey does not provide precise annual statistics for the whole population and hence much less so for subpopulations. This problem is likely more severe in most surveys, which typically collect smaller samples.

I use the internal ACS file that includes group quarters and use sampling weights throughout, so the sample is representative of the entire residential population in NY. Thus, the recipients in my data are a sample from all SNAP recipients, except for those who neither live in housing units nor group quarters. They are a small fraction of recipients, such as homeless recipients not living in temporary shelters.⁸ In the absence of non-sampling error, the recipients in the survey would therefore be very close to a random sample from the population of recipients in the administrative SNAP data. Consequently, one can compare the survey estimates to the totals from the administrative micro data to evaluate the estimators.

In order to further improve comparability of the survey estimates and the totals from the administrative data, I link the administrative records to the survey data at the individual level. I then replace the survey response on SNAP receipt by the administrative measure. While one would usually have to rely on survey reports, using the linked administrative measure makes the evaluation cleaner by abstracting from sources of non-sampling error such as misreporting or non-response, which often loom large in survey data on income and transfer programs (Meyer and Mittag, 2018). The ability of the estimators to address such non-sampling error depends on the nature of this error and hence the survey at hand. Therefore, drawing general conclusions about their performance in the presence of non-sampling error is difficult. Linking the data provides me with the exact same measure for

⁸I exclude them from the administrative data as far as possible, but the population excluded from the ACS is not cleanly identified in the administrative data. Both the share of the individuals I exclude and those not covered by the ACS by design are too small to affect any conclusions substantively.

the ACS and the entire population and thereby provides a cleaner comparison of the ability of the estimators to reduce sampling variation.

The ACS data are linked to the administrative records described above using the Person Identification Validation System (PVS) of the U.S. Census Bureau. See Wagner and Layne (2012) for a detailed description. In short, the PVS uses the person data (such as address, name, gender, and date of birth) from the administrative records and survey data to find a matching record in a reference file that contains all transactions recorded against a social security number. If a matching record is found, the transformed social security number (PIK) of the record from the reference file is attached to the corresponding records in the data. A PIK is obtained for over 99 percent of the administrative records and 89 percent of individuals in the ACS. I adjust the ACS sampling weights using inverse probability weighting (Wooldridge, 2007) to address the issue of missing PIKs in the survey.⁹ The coefficients of the Probit model I use to predict these probabilities are reported in Appendix Table A1.

The accuracy of the linked data appears to be high, as further discussed in Celhay, Meyer and Mittag (2018*a,b*) and Meyer and Mittag (forthcoming). Using the linked administrative measure instead of survey reports addresses both misreporting and non-response, since reciprocity status is accurately recorded for all linked individuals. The administrative measure is also not affected by post-processing and imputation. Yet, it may still differ from administrative totals due to coverage and linkage errors. Thus, the administrative measure in the linked data is not free of error, but it should isolate sampling error as much as possible. The previous literature (see e.g. Meyer, Mittag and Goerge, 2018) provides evidence that errors from imperfect linkage are likely to be small compared to errors in survey reports. Nonetheless, the estimate of the total number of recipients in NY from the linked data falls short of the total number in the administrative data by 11 percent. This difference is likely to be a combination of sampling variation at the state level, failure to link some individuals (e.g. those moved out of state) and undercoverage of the ACS (Meyer and Mittag, 2018).

⁹Only few PIKs are missing in the administrative data, so I do not attempt to adjust for this problem.

Most survey data suffer from similar or worse problems and the estimators, particularly the population adjusted estimator may improve this problem. Thus, I compare unadjusted results in order to assess how well the estimators work with contaminated data. To focus on the ability of the estimators to reduce sampling variation, I also scale up the county level estimates so that their annual sum matches the state total. This benchmarking reduces the non-sampling error due to the differences in the data for the entire state.

The population adjusted estimator additionally uses aggregate information on total receipt in the overall population, which is often known even for small areas. For SNAP, annual total receipt by county is published by the US Department of Agriculture (USDA). In the application below, I use total receipt from the administrative micro data instead of the BEA numbers to make definitions more comparable.¹⁰ As is often the case, the USDA provides statistics for the overall population, but not for any subpopulations. Thus, researchers and policy makers interested in receipt rates of subpopulations, such as the elderly, have to rely on other sources, such as survey data.

4 Results

I use the estimators above and the 2008-2012 linked ACS data to obtain annual estimates of the total number of female and elderly recipients for 61 counties in NY.¹¹ Table 1 reports measures of how these estimates differ from the true number of female and elderly recipients according to the administrative micro data.¹² To put the magnitudes in context, the first row provides the average number of recipients per county according to each estimator. The second row reports the percentage difference between the estimated and the administrative total for the entire state as a measure of how survey and administrative data differ overall. The

¹⁰The official statistics are similar, but report average monthly receipt rather than the more common survey definition of receipt at any time in the past 12 months.

¹¹The ACS sample for Hamilton county is too small to allow disclosure of county-level estimates, so I pool Herkimer and Hamilton county throughout, making it 61 estimates for 62 counties. For the elderly, sample sizes for Seneca and Orleans prevented disclosure of the county totals for 2010. While these counties were included in the estimation, they are excluded from the results for the elderly reported here.

¹²Full tables with estimates by county by year are available in a web appendix.

remaining rows contain root mean square error (RMSE) and mean absolute deviation (MAD) as measures of county-level error in the estimates. RMSE is defined as $\sqrt{N^{-1} \sum (\hat{y}_{ct} - y_{ct})^2}$ and MAD as $N^{-1} \sum |\hat{y}_{ct} - y_{ct}|$, where N denotes the number of county-year estimates and the sum is taken over these estimates. The table also provides these measures in percent of the error of the direct survey estimates for comparison.

In order to abstract from differences between the administrative and the survey data for the entire state, the lower part of Table 1 scales county-year estimates up to match the number of recipients in the subpopulation for the entire state.¹³ Comparing the top and bottom panels of table 1 shows that this simple benchmarking to state totals substantially improves accuracy. This rescaling is only feasible if the statistic of interest is available for the subpopulation in a larger area. To be unbiased, it requires the survey errors to be randomly distributed across counties. Both the source of errors and the results suggest that this assumption is a reasonable approximation here. Nevertheless, this condition and hence that scaling up results reduces error may not generalize.

Since survey totals often differ from the true totals for reasons other than sampling variation, the unscaled results arguably resemble a typical application more closely and allow me to evaluate the potential of the estimators to address non-sampling error. Even so, non-sampling error likely differs between surveys, so it is not clear whether these findings generalize. The scaled up results isolate sampling error better, so they are more indicative of the capacity of the estimators to reduce it. Sampling error does not differ qualitatively between surveys, so findings from the rescaled results should be more likely to generalize.

The second row of table 1 confirms that the population adjusted estimator in column 3 and 6 can reduce bias, but will not always do so. The results for females in column 3 show that the population adjusted estimator can reduce non-sampling error, as it cuts bias by 70 percent. For the elderly, bias slightly increases compared to the direct survey estimator. In both cases the difference switches signs, i.e. contrary to the direct survey estimator, the

¹³This benchmarking method is very simple, see Pfeiffermann (2013) for a discussion of more sophisticated benchmarking methods, which would likely reduce the variance of all three estimators.

Table 1: Comparing Estimators for Subpopulations - Female and Elderly SNAP Recipients

Estimator:	<i>Females</i>			<i>Elderly</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
	Direct Survey	Fixed Effects	Pop. Adjusted	Direct Survey	Fixed Effects	Pop. Adjusted
Av. Recipients per County	27,144	27,238	30,578	7,349	7,417	8,335
Difference to Admin. Total	-8.9%	-8.6%	2.6%	-5.7%	-4.8%	6.9%
Root Mean Squared Error	8,168	6,916	2,711	2,298	1,921	1,694
... in % of Direct Est.		84.7%	33.2%		83.6%	73.7%
Mean Absolute Deviation	3,211	2,960	1,061	862	640	778
... in % of Direct Est.		92.2%	33.0%		74.3%	90.2%
<i>Rescaled to Match State Total</i>						
Root Mean Squared Error	5,983	4,930	1,371	1,465	776	1,051
... in % of Direct Est.		82.4%	22.9%		53.0%	71.7%
Mean Absolute Deviation	656	354	542	2,339	1,754	701
... in % of Direct Est.		75.0%	30.0%		54.0%	82.6%

Note: All statistics based on annual total recipients per county from the 2008-2012 ACS linked to NY OTDA records. Elderly are those 60 and over. RMSE and MAD calculated from differences to county totals from the administrative microdata. For the rescaled statistics, the annual county totals are multiplied by the ratio of the true to the estimated state total for the corresponding year and estimator.

population adjusted estimator overstates the number of recipients. As section 2 shows, this implies that the survey captures both subpopulations better than the general population. The fixed effects estimator in column 2 and 5 does not significantly affect the difference to state totals.¹⁴ The finding that the fixed effects estimator does not improve bias from non-sampling sources is likely to hold more generally, because this estimator does not incorporate any information besides the survey data.

The remaining rows of table 1 show that both estimators substantially reduce error. They improve both measures of error in the unscaled and rescaled results for both subpopulations. For females, the fixed effects estimator reduces RMSE and MAD by 8 to 25 percent compared to the direct survey estimator. Error reductions are much larger for the population adjusted estimator at 67 to 77 percent. For the elderly, the fixed effects estimator yields larger error reductions, sometimes cutting average error almost in half. Here, it performs better or

¹⁴Aggregate totals for the fixed effects estimator only differ from those implied by the survey because I use population counts by county from official sources rather than those implied by the survey data.

on par with the population adjusted estimator, which is in line with the arguments from section 2 that the fixed effects estimator should yield larger error reductions compared to the alternatives when the annual samples are particularly small. The larger bias of the population adjusted estimator for the elderly may further contribute to the better relative performance of the fixed effects estimator for the elderly.

To further examine the reasons for the relative difference in performance and thereby the conditions under which each estimator is preferable, table 2 reports the same measures of error separately by county size. The upper panel contains error measures for the estimates where the population of interest is below the median populations size, the lower panel reports the same measures for the estimates with above median population sizes. For females, the population adjusted estimator still performs better than the fixed effects estimator regardless of population size. For the elderly, the fixed effects estimator still performs better. The population adjusted estimator actually performs worse than the direct survey estimator for the elderly for small populations when the results are not rescaled. The error reduction from rescaling shows that the performance of the population adjusted estimator is affected by the bias, but this is not the main cause of the difference in performance. Rather, comparing the error reductions relative to the direct survey estimate shows that the error reductions of the fixed effects estimator are larger for the *smaller* counties in the top panel of table 2 in all cases. For the population adjusted estimator, error reductions are larger for the *larger* counties in the bottom panel of table 2 in all cases.

Consequently, these results confirm that the fixed effects estimator is attractive when the samples are particularly small and the main concern is sampling variation rather than other sources of survey error. As discussed in section 2, one may expect this advantage to hold more generally, simply because additional observations reduce variance more in small samples. Even so, the fixed effects estimator requires repeated cross-sections of the survey. If there are county-specific time effects (which is testable), it may be biased or even perform worse than the direct survey estimator. The population adjusted estimator can yield larger error

Table 2: Comparing Estimators by Population Size

Estimator:	<i>Females</i>			<i>Elderly</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
	Direct Svy	Fixed Effects	Pop. Adjusted	Direct Svy	Fixed Effects	Pop. Adjusted
Smaller Counties						
Av. Recipients per County	3,552	3,565	4,170	755	761	897
Difference to Admin. Total	-13%	-13%	2%	-8%	-7%	10%
Root Mean Squared Error	940	679	363	264	127	299
...in % of Direct Est.		72%	39%		48%	113%
Mean Absolute Deviation	716	563	286	215	100	230
...in % of Direct Est.		79%	40%		47%	107%
<i>Rescaled to Match State Total</i>						
Root Mean Squared Error	844	503	345	270	118	270
...in % of Direct Est.		60%	41%		44%	100%
Mean Absolute Deviation	659	388	276	217	91	208
...in % of Direct Est.		59%	42%		42%	96%
Larger Counties						
Av. Recipients per County	51,523	51,701	57,866	14,164	14,297	16,024
Difference to Admin. Total	-9%	-8%	3%	-6%	-5%	7%
Root Mean Squared Error	11,609	9,837	3,848	3,265	2,737	2,397
...in % of Direct Est.		85%	33%		84%	73%
Mean Absolute Deviation	5,790	5,438	1,862	1,531	1,199	1,345
...in % of Direct Est.		94%	32%		78%	88%
<i>Rescaled to Match State Total</i>						
Root Mean Squared Error	8,488	7,011	1,923	2,071	1,100	1,473
...in % of Direct Est.		83%	23%		53%	71%
Mean Absolute Deviation	4,076	3,166	1,140	1,110	627	887
...in % of Direct Est.		78%	28%		56%	80%

Note: All statistics based on annual total recipients per county from the 2008-2012 ACS linked to NY OTDA records. Elderly are those 60 and over. County-year estimates are split into a smaller group (top panel) and a larger group (bottom panel) by the size of the population of interest (according to the administrative data). RMSE and MAD calculated from differences to county totals from the administrative microdata. For the rescaled statistics, the annual county totals are multiplied by the ratio of the true to the estimated state total for the corresponding year and estimator.

reductions when the statistic of interest is known for the overall population. Its advantage vanishes for particularly small populations, which likely generalizes to other cases in which estimation noise is dominated by the error in f , as discussed in section 2. The advantage is pronounced for females, where the estimator also substantially reduces bias, which suggests that the population adjusted estimator performs particularly well when survey errors other than sampling error are a concern. The estimator improves such non-sampling error when the survey coverage rates of the overall and the subpopulation are similar (as defined in section 2). Such improvements likely arise more generally when survey errors do not depend on subpopulation status, but the population adjusted estimator can perform worse than the direct survey estimates when survey error differs between the subpopulation and the overall population in specific ways.

Overall, the empirical analysis underlines that even simple data combination methods can improve precision, sometimes dramatically so, compared to the commonly used direct survey estimates. Both estimators can yield large error reductions and the results provide evidence on the conditions under which each estimator performs well. These conditions are straightforward and can often be assessed empirically. More advanced methods (e.g. Datta et al., 1999, You, 2008) and more sophisticated applications, such as the ones by the U.S. Census Bureau (Bauder, Luery and Szelepka, 2018, Luery, 2010) and Statistics Netherlands (van den Brakel and Krieg, 2016, Bollineni-Balabay, van den Brakel and Palm, 2016), provide ample ways to relax assumptions or further increase accuracy.

5 Conclusions

This report discusses and evaluates two simple methods to estimate official statistics for small subpopulations. I use administrative SNAP records and administrative records linked to survey data to evaluate these methods, but implementing them in practice only requires survey data and aggregate statistics. I find that both estimators improve over the commonly

used direct survey estimates in an application to SNAP receipt.

In combination with known properties of the estimators, the results suggest conditions under which the improvements generalize. Thereby, they can guide practitioners in choosing an estimator. The fixed effects estimator is especially attractive when the samples are particularly small. It performs well when the main concern is sampling variation rather than other types of survey error. The population adjusted estimator yields large error reductions, making it appealing when the statistic of interest is available for the entire population. It becomes less appealing relative to the fixed effects estimator as the population of interest becomes smaller. Contrary to the fixed effects estimator, the population adjusted estimator can also mitigate other sources of survey error, such as misreporting or non-response, but it can also increase error. Improvements are likely if the errors are similar in the subpopulation of interest and the overall population. This advantage will be particularly pronounced if the coverage rates of the overall and the subpopulation are similar, but differ from one. Section 2 provides precise conditions that can often be assessed empirically.

The literature provides many potential ways to relax assumptions or better model the dependence of the effects across space or time as discussed in section 2. For example, one could allow the time-specific effects to depend on covariates such as local unemployment rates or model their time-series properties. Both estimators could also be improved by pooling areas or time periods optimally, potentially using machine learning methods as in Bonhomme and Manresa (2015) or Athey and Imbens (2016). For the fixed effects estimator, one could pool areas or time periods for which the fixed effects do not differ. For the population adjusted estimator, accuracy could be improved further by estimating f_{ct} from the largest samples for which the fraction of recipients in the subpopulation is constant. One could also use these methods to find county groups with common time specific effects to relax the assumption of the fixed effects estimator that these shocks are common to all areas.

In conclusion, even these two simple data combination estimators can yield sizable error reductions. The conditions under which these two estimators or more complex methods are

desirable depend on the application at hand. Both the theoretical and the empirical analyses emphasize conditions under which each estimator performs well that can be assessed empirically. Thereby, they provide practical guidance when choosing an estimator and underline that it is feasible to improve the accuracy of direct survey estimates of official statistics for small subpopulations in practice.

Appendix Table A1: The Determinants of an Individual Having a PIK, Probit Coefficients

	2008		2009		2010		2011		2012	
	<i>Coef.</i>	<i>SE</i>	<i>Coef.</i>	<i>SE</i>	<i>Coef.</i>	<i>SE</i>	<i>Coef.</i>	<i>SE</i>	<i>Coef.</i>	<i>SE</i>
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
HH Type: Multiple Adults	-0.012	0.012	-0.046	0.012	-0.001	0.012	-0.033	0.012	0.034	0.012
Non-Farm Household	-0.114	0.016	-0.183	0.016	-0.168	0.017	-0.291	0.016	-0.250	0.015
Group Quarter	-0.080	0.050	-0.115	0.041	0.090	0.049	-0.802	0.042	-0.902	0.041
Child Under 6 Present	0.007	0.016	-0.012	0.015	0.060	0.017	0.004	0.016	-0.010	0.016
Child Aged 6-17 Present	0.091	0.012	0.136	0.011	0.077	0.012	0.062	0.012	0.071	0.012
Age 16-29	-0.051	0.014	-0.106	0.013	-0.055	0.014	-0.118	0.014	-0.068	0.014
Age 30-39	-0.051	0.016	-0.070	0.016	-0.045	0.017	-0.059	0.017	-0.064	0.016
Age 50-59	0.095	0.017	0.100	0.016	0.049	0.017	0.081	0.017	0.043	0.016
Age 60-69	0.150	0.019	0.152	0.019	0.053	0.020	0.112	0.020	0.094	0.019
Age 70 or More	0.173	0.020	0.233	0.020	0.085	0.021	0.132	0.022	0.130	0.020
HHlder: Less Than HS	-0.075	0.017	-0.049	0.016	-0.072	0.017	-0.017	0.017	-0.053	0.016
HHlder: HS Graduate	-0.064	0.015	-0.074	0.014	-0.043	0.015	-0.070	0.015	-0.054	0.014
HHlder: Graduate Degree	-0.025	0.016	0.005	0.015	-0.045	0.016	-0.021	0.016	-0.038	0.015
Less Than HS	-0.021	0.017	-0.050	0.016	-0.035	0.018	-0.059	0.017	0.000	0.016
HS Graduate	-0.060	0.017	-0.067	0.016	-0.068	0.017	-0.089	0.017	-0.046	0.016
Graduate Degree	0.033	0.018	0.008	0.018	0.039	0.019	-0.014	0.019	0.036	0.018
Currently in School/OOU	-0.035	0.029	-0.482	0.026	-0.121	0.030	-0.277	0.028	-0.069	0.029
Race: White	-0.074	0.032	-0.055	0.032	-0.148	0.031	0.006	0.028	-0.003	0.027
Race: Black	-0.105	0.034	-0.090	0.033	-0.128	0.033	-0.034	0.029	0.025	0.028
Race: AI or AN	-0.312	0.074	-0.204	0.076	-0.280	0.074	-0.173	0.070	0.002	0.066
Race: Asian	-0.115	0.036	-0.152	0.036	-0.224	0.035	-0.084	0.032	-0.047	0.031
Race: Other	0.038	0.036	0.044	0.035	-0.055	0.035	0.003	0.031	-0.014	0.030
Hispanic	0.015	0.016	-0.032	0.016	0.017	0.017	0.034	0.017	0.028	0.016
Not Working This Year	-0.221	0.015	-0.176	0.014	-0.160	0.015	-0.138	0.015	-0.165	0.014
Full-Time Employed	-0.076	0.014	-0.048	0.014	-0.100	0.015	-0.082	0.015	-0.087	0.014
HHlder: Non-Citizen	-0.205	0.017	-0.185	0.017	-0.124	0.017	-0.143	0.017	-0.268	0.016
Non-Citizen	-0.511	0.018	-0.581	0.018	-0.396	0.019	-0.440	0.018	-0.339	0.017
HH in Poverty	-0.059	0.015	-0.077	0.014	-0.058	0.015	-0.083	0.014	-0.047	0.014
Income 100 - 130% of FPL	0.029	0.021	-0.020	0.020	-0.034	0.021	-0.021	0.020	0.029	0.020
Income 130 - 200% of FPL	-0.105	0.014	-0.059	0.014	-0.025	0.015	-0.020	0.015	-0.002	0.014
Poverty Status Missing	-0.012	0.051	0.032	0.045	-0.050	0.050	-0.090	0.046	0.099	0.045
Only English Spoken in HH	0.138	0.012	0.080	0.012	0.102	0.013	0.084	0.013	0.106	0.012
Speaks Poor English	-0.315	0.018	-0.323	0.018	-0.265	0.018	-0.248	0.018	-0.199	0.017
Disabled Present in HH	0.049	0.014	0.045	0.014	0.067	0.015	0.030	0.015	0.049	0.014
Disabled	0.124	0.020	0.141	0.020	0.144	0.022	0.150	0.021	0.147	0.020
Rural	-0.015	0.013	0.017	0.013	0.043	0.015	0.060	0.014	0.074	0.014
Interview Mode: Mail	0.911	0.010	0.902	0.010	0.955	0.011	0.911	0.011	0.931	0.010
Interview Mode: CATI	-0.131	0.012	-0.032	0.011	-0.150	0.012	-0.032	0.012	-0.065	0.012
Constant	1.170	0.040	1.153	0.039	1.362	0.039	1.276	0.036	1.159	0.034
Number of Observations	276,146		276,304		276,002		317,085		337,161	

Notes: All analyses use individual weights. The omitted categories are: Single-Adult HH (HH type), 40-49 (Age), Some College (Education), Multiple (Race), Part-time (Employment), >200% FPL (Income), CATI (Interview Mode).

References

- Athey, Susan, and Guido Imbens.** 2016. “Recursive partitioning for heterogeneous causal effects.” *Proceedings of the National Academy of Sciences*, 113(27): 7353–7360.
- Bauder, Mark, Donald Luery, and Sam Szelepka.** 2018. “Small Area Estimation of Health Insurance Coverage in 2010-2016.” U.S. Census Bureau Technical Report, Washington, D.C.
- Bohn, Sarah, Caroline Danielson, Matt Levin, Marybeth Mattingly, and Christopher Wimer.** 2013. “The California poverty measure: A new look at the social safety net.” Public Policy Institute of California Policy Report.
- Bollineni-Balabay, Oksana, Jan van den Brakel, and Franz Palm.** 2016. “Multivariate state space approach to variance reduction in series with level and variance breaks due to survey redesigns.” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 179(2): 377–402.
- Bonhomme, Stephane, and Elena Manresa.** 2015. “Grouped Patterns of Heterogeneity in Panel Data.” *Econometrica*, 83(3): 1147–1184.
- Celhay, Pablo, Bruce D. Meyer, and Nikolas Mittag.** 2018*a*. “Errors in Reporting and Imputation of Government Benefits and Their Implications.” Unpublished Manuscript.
- Celhay, Pablo, Bruce D. Meyer, and Nikolas Mittag.** 2018*b*. “What Leads to Measurement Errors? Evidence from Reports of Program Participation from Multiple Surveys.” Unpublished Manuscript.
- Cerf Harris, Benjamin.** 2014. “Within and Across County Variation in SNAP Misreporting: Evidence from Linked ACS and Administrative Records.” U.S. Census Bureau CARRA Working Paper No. 2014-05.
- Datta, Gauri S., Partha Lahiri, Tapabrata Maiti, and Kim L. Lu.** 1999. “Hierarchical Bayes estimation of unemployment rates for the states of the US.” *Journal of the American Statistical Association*, 94(448): 1074–1082.
- Deville, Jean-Claude, and Carl-Erik Särndal.** 1992. “Calibration estimators in survey sampling.” *Journal of the American Statistical Association*, 87(418): 376–382.
- D’Onofrio, Christine, John Krampner, Danny Silitonga, Jihyun Shin, and Vicky Virgin.** 2015. “The CEO Poverty Measure, 2005-2013.” Center for Economic Opportunity Policy Report.
- Gee, Geof, and Robin Fisher.** 2004. “Errors-In-Variables County Poverty and Income Models.” American Statistical Association Proceedings of the Section on Government and Social Statistics.
- Haider, Steven J., Alison Jacknowitz, and Robert F. Schoeni.** 2003. “Food Stamps and the Elderly: Why Is Participation so Low?” *The Journal of Human Resources*, 38: 1080–1111.

- Harter, Rachel, Michael Macaluso, and Kirk Wolter.** 1999. "Small domain estimation of employment using CES and ES202 data." Federal Committee on Statistical Methodology Research Conference Proceedings.
- Harter, Rachel, Michael Macaluso, and Kirk Wolter.** 2005. "Evaluating the Fundamentals of a Small Domain Estimator." *Survey Methodology*, 29(1): 63–70.
- Hoynes, Hilary W, and Diane Whitmore Schanzenbach.** 2016. "US Food and Nutrition Programs." In *Economics of Means-Tested Transfer Programs in the United States*. Vol. 1, , ed. Robert A. Moffitt, Chapter 3, 219–302. Chicago:University of Chicago Press.
- Imbens, Guido W., and Tony Lancaster.** 1994. "Combining Micro and Macro Data in Microeconomic Models." *The Review of Economic Studies*, 61(4): 655–680.
- Kim, Jae Kwang, Seunghwan Park, and Seo-young Kim.** 2015. "Small area estimation combining information from several sources." *Survey Methodology*, 41(1): 21–36.
- Kott, Phillip S.** 2009. "Calibration weighting: Combining probability samples and linear prediction models." In *Handbook of Statistics 29B, Sample Surveys: Inference and Analysis*, ed. Danny Pfeffermann and Calyampudi R. Rao, 55–82. Amsterdam:North Holland.
- Little, Roderick J.A.** 2012. "Calibrated Bayes, an Alternative Inferential Paradigm for Official Statistics." *Journal of Official Statistics*, 28(3): 309–334.
- Luery, Donald.** 2010. "Small Area Income and Poverty Estimates Program." U.S. Census Bureau Technical Report, Washington, D.C.
- Meyer, Bruce D., and Nikolas Mittag.** 2018. "An Empirical Total Survey Error Decomposition Using Data Combination." Unpublished Manuscript.
- Meyer, Bruce D., and Nikolas Mittag.** forthcoming. "Using Linked Survey and Administrative Data to Better Measure Income: Implications for Poverty, Program Effectiveness and Holes in the Safety Net." *American Economic Journal: Applied Economics*.
- Meyer, Bruce D., Nikolas Mittag, and Robert M. Goerge.** 2018. "Errors in Survey Reporting and Imputation and Their Effects on Estimates of Food Stamp Program Participation." Unpublished Manuscript.
- Pfeffermann, Danny.** 2002. "Small Area Estimation: New Developments and Directions." *International Statistical Review / Revue Internationale de Statistique*, 70(1): 125–143.
- Pfeffermann, Danny.** 2013. "New Important Developments in Small Area Estimation." *Statistical Science*, 28(1): 40–68.
- Pfeffermann, Danny, and L. Burck.** 1990. "Robust small area estimation combining time series and cross-sectional data." *Survey Methodology*, 16: 217–237.
- Pfeffermann, Danny, and Richard Tiller.** 2006. "Small-area estimation with state–space models subject to benchmark constraints." *Journal of the American Statistical Association*, 101(476): 1387–1397.

- Pfeffermann, Danny, and Susana R. Bleuer.** 1993. "Robust joint modelling of labour force series of small areas." *Survey Methodology*, 19(2): 149–163.
- Rao, Jon N. K., and Isabel Molina.** 2015. *Small Area Estimation*. . 2nd ed., New York:Wiley.
- Rao, Jon N. K., and Mingyu Yu.** 1994. "Small-area estimation by combining time-series and cross-sectional data." *Canadian Journal of Statistics*, 22(4): 511–528.
- Scherpf, Erik, Constance Newman, and Mark Prell.** 2014. "Targeting of Supplemental Nutrition Assistance Program Benefits: Evidence from ACS and NY SNAP Administrative Records." USDA Working Paper.
- Tarozzi, Alessandro, and Angus Deaton.** 2009. "Using Census and Survey Data to Estimate Poverty and Inequality for Small Areas." *Review of Economics and Statistics*, 91(4): 773–792.
- Tiller, Richard B.** 1992. "Time series modeling of sample survey data from the US Current Population Survey." *Journal of Official Statistics*, 8(2): 149–166.
- U.S. Census Bureau.** 2008. "A Compass for Understanding and Using American Community Survey Data: What General Data Users Need to Know." U.S. Government Printing Office, Washington, D.C.
- U.S. Census Bureau.** 2014. "American Community Survey: Design and Methodology." U.S. Census Bureau Technical Paper, Washington, D.C.
- U.S. Census Bureau.** 2015. "Income and Poverty in the United States: 2014." U.S. Census Bureau Current Population Report P60-252, Washington, D.C.
- van den Brakel, Jan A., and Sabine Krieg.** 2016. "Small area estimation with state space common factor models for rotating panels." *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 179(3): 763–791.
- Wagner, Deborah, and Mary Layne.** 2012. "The Person Identification Validation System (PVS): Applying the Center for Administrative Records Research & Applications' (CARRA) Record Linkage Software." U.S. Census Bureau Center for Administrative Records Research and Applications Internal Paper.
- West, Sandra A.** 1983. "A comparison of different ratio and regression type estimators for the total of a finite population." American Statistical Association Proceedings of the Section on Survey Research Methods.
- Wooldridge, Jeffrey M.** 2007. "Inverse probability weighted estimation for general missing data problems." *Journal of Econometrics*, 141(2): 1281–1301.
- You, Yong.** 2008. "An integrated modeling approach to unemployment rate estimation for sub-provincial areas of Canada." *Survey Methodology*, 34(1): 19–27.

- You, Yong, Jon N.K. Rao, and Jack Gambino.** 2003. “Model-based unemployment rate estimation for the Canadian Labour Force Survey: a hierarchical Bayes approach.” *Survey Methodology*, 29(1): 25–32.
- Zanutto, Elaine, and Alan Zaslavsky.** 2002. “Using administrative records to improve small area estimation: An example from the US Decennial Census.” *Journal of Official Statistics*, 18(4): 559.