

REDESIGN OF THE MEDICARE CURRENT BENEFICIARY SURVEY SAMPLE

**Annie Lo, Adam Chu, and Richard Apodaca, Westat
Annie Lo, Westat, 1650 Research Boulevard, Rockville, Maryland 20850**

Key Words: Multistage probability sample, resampling methods, Ernst method for overlap maximization, longitudinal sample, rotating panels, Medicare population

1. Introduction

The Medicare Current Beneficiary Survey (MCBS) conducted by the Centers for Medicare and Medicaid Services (CMS), Department of Health and Human Services, is a continuous sample survey of Medicare beneficiaries residing in the United States and Puerto Rico. The MCBS collects data on access to health care, health status, source of care, health care utilization and costs, satisfaction with health care, and other health-related topics (e.g., see Sharma, Chan, Liu, and Ginsberg, 2001). A representative sample of Medicare beneficiaries (referred to as a “panel”) is selected for the MCBS each year using a stratified multistage probability sample design. The sample of first-stage or primary sampling units (PSUs), which includes MSAs (metropolitan statistical areas) and groups of rural (nonMSA) counties, was designed and selected in 1991. Although new beneficiary samples are selected each year to supplement the original sample, the new samples are always selected from the same PSUs. Over time, the continued use of the original PSU sample has resulted in losses in both sampling precision and operational efficiency. In 2000, based on an evaluation of the existing PSU sample, a decision was made to reselect the PSUs. This paper summarizes some of the analyses leading to that decision and describes the procedures used to update and select the new MCBS PSU sample.

2. The MCBS Sample Design

The MCBS employs a stratified multistage probability sample with three stages of selection. The first stage involved the selection of PSUs consisting of MSAs and groups of rural counties. The PSUs were selected with probabilities proportionate to 1980 population within strata defined by Census region, metropolitan status, and selected PSU-level socio-economic characteristics. Two PSUs were selected per stratum. The second sampling stage consisted of the selection of ZIP Code areas within each sampled PSU. To facilitate linking with available county-level data, the second-stage sampling unit was defined to be the part of the ZIP Code area that was physically contained within a given county. In other words, ZIP Code areas that crossed county borders were subdivided by county into separate units called “ZIP fragments.” For sampling purposes, small

ZIP fragments were combined into clusters where necessary to ensure that each ZIP cluster would provide a reasonable workload for interviewers if selected for the sample. At the third and final stage of selection, beneficiaries within the sampled ZIP clusters were stratified by age and subsampled at rates designed to yield self-weighting (equal probability) samples of beneficiaries within each of seven age groups. Additional details about the original MCBS sample design are provided in Apodaca, Judkins, Lo, and Skellan (1992).

The MCBS was originally intended to be a true longitudinal survey in which sampled Medicare beneficiaries would be interviewed three times a year throughout the remainder of their lives. However, after two years of data collection, it became clear that this would be impractical. Thus, a decision was made to switch from a fixed panel design to a rotating panel design in which roughly one-third of the existing sample (i.e., the oldest panel) is retired each year, and a new panel is selected to replace it. Under this design, beneficiaries in each newly selected panel are interviewed three times a year for a maximum of four years. Table 1 illustrates the basic features of the rotating panel design developed for the MCBS. Additional details are given in Westat (2001).

Table 1. Panel rotation scheme for the MCBS*

Panel year	Data collection year				
	1998	1999	2000	2001	2002
1994	A4	—	—	—	—
1995	B3	B4	—	—	—
1996	C2	C3	C4	—	—
1997	D1	D2	D3	D4	—
1998	—	E1	E2	E3	E4
1999	—	—	F1	F2	F3
2000	—	—	—	G1	G2
2001	—	—	—	—	H1

*Panel year refers to the year of the fall round in which the panel is introduced into the study. Data collection year refers to subsequent data collection rounds. The letters A, B, C, etc. are used to designate a particular panel. The numeric values indicate the data collection year. For example, C4 refers to the fourth year of data collection for the 1996 panel.

Over 15,000 beneficiaries were selected for the initial round of the MCBS. In each of the following two years, supplemental samples of about 2,400 beneficiaries per year were added to the original sample to compensate for sample attrition and to give coverage to newly enrolled Medicare beneficiaries. With the implementation of the rotating panel design in 1994, the number of beneficiaries

selected for each annual supplement (i.e., nationally representative panel) has been between 6,300 to 6,400 beneficiaries per year.

3. Considerations for Updating MCBS PSUs

The main reasons for updating and reselecting an existing PSU sample are: (a) to improve sampling precision through the use of more up-to-date sampling measures of size and stratification schemes, and (b) to maintain more balanced sample workloads across PSUs. Kish (1965, page 482) also cites the need to avoid “inertia in continuing operations” that can hinder improvements of outdated and inefficient procedures. The Current Population Survey (CPS), for example, has traditionally updated its PSU sample at 10-year intervals using data from the most recent decennial Census to redesign the sample (U.S. Department of Labor, 2000).

For the MCBS, there was evidence that design effects (defined to be ratio of the variance of an estimate derived from the MCBS to the corresponding variance based on a simple random sample of the same size) were increasing over time. For a number of statistics examined in Westat (2000), the median design effect for the total beneficiary sample increased by an average of four percent annually between 1992 and 1995. By age group, the average increase in median design effects varied from less than two percent for older beneficiaries (75 years or older) to around three to six percent annually for younger beneficiaries (74 years or younger). Although the estimated design effects fluctuated widely from year to year, the overall patterns did suggest that design effects had generally increased between 1992 and 1995. By 1996, however, there was a noticeable drop in design effects, probably due to the fact that over two-thirds of the original MCBS panel had been phased out of the study by this time. While design effects after 1996 did not increase as greatly as in previous years, there did appear to be modest increases for some subgroups. The results shown in Table 2 for selected health related variables illustrate the magnitude and variability of the change in design effects over time.

Table 2. MCBS design effects, 1992-1999

Survey year	Characteristic				Hyper-tension
	Functional limitation*				
	None	IADL only	1 or 2 ADLs	3 to 5 ADLs	
1992	2.00	1.29	1.28	1.44	1.55
1993	1.78	1.46	1.45	1.07	1.52
1994	1.91	1.34	1.52	0.55	1.61
1995	1.75	1.38	1.04	1.06	1.50
1996	1.30	1.38	1.29	1.36	1.71
1997	1.33	1.29	1.17	1.45	1.82
1998	1.30	1.41	1.10	1.08	1.86
1999	1.84	1.35	1.97	1.00	NA

*Instrumental activities of daily living (IADLs). Activities of daily living (ADLs).

It should be noted that the design effects summarized in Table 2 reflect the increase in variance

arising from a variety of sources. The continued use of an increasingly inefficient PSU sampling measure of size can lead to both increased clustering effects as well as increased variation in sampling weights. In Section 3.1, the variance of an estimate based on the MCBS design is decomposed to show how various design features affect overall sampling precision.

In addition to increased variances, the use of an old PSU sample can lead to a less efficient distribution of workload across PSUs. This occurs because the measure of size used to select the original PSU sample may no longer adequately control PSU workloads (sample sizes). For the MCBS, the original PSU measure of size was 1980 population. Not only was the source of data 10 years older than the PSU sample, the measure of size assigned to PSUs reflected total U.S. population rather than the Medicare population. Updating the PSU measure of size with current counts of Medicare beneficiaries, therefore, was expected to ameliorate the worsening imbalance in the PSU workloads. Implications of the aging PSU sample on survey operations and costs are discussed further in Section 3.2.

3.1 Precision of Estimates

As is the case with virtually all sample surveys, approximately unbiased estimates of totals derived from the MCBS are weighted sums of the form:

$$\hat{y} = \sum_{p=1}^P \sum_{i=1}^{n_p} w_{pi} y_{pi} , \tag{1}$$

where y_{pi} is the observed value of the characteristic being estimated for the i -th sample beneficiary in panel p , and w_{pi} is the corresponding sampling weight. Annual cost and use estimates are typically based on three complete (continuing) panels, while access-to-care estimates are based on four panels. The panel sample sizes, n_p , vary slightly with newer panels being somewhat larger than older ones due to attrition. The sampling weights, w_{pi} , reflect the beneficiaries’ overall probabilities of selection, and include adjustments for nonresponse and undercoverage. Additional details about the weighting procedures employed in the MCBS are given in Judkins and Lo (1993).

To bring out important features of the sample design employed for the MCBS, it is useful to express the estimated total given by equation (1) in the following alternative form:

$$\hat{y} = \sum_{g=1}^G \hat{y}_g = \sum_{g=1}^G \sum_{p=1}^P a_{gp} \hat{y}_{gp} , \tag{2}$$

where

$$\hat{y}_{gp} = \sum_{i=1}^{n_{gp}} w_{gpi}^{NR} y_{gpi} \tag{3}$$

is the estimated total for the g -th “combination group” in panel p . A combination group is a subset of the beneficiary sample within which the individual panel-specific estimates are “composited” to form an overall combined estimate. Note that the weights, w_{gpi}^{NR} , in equation (3) are panel-specific nonresponse-adjusted weights that inflate the results for panel p to population levels; thus,

$$\hat{y}_g = \sum_{p=1}^P a_{gp} \hat{y}_{gp}$$

is a composite estimate of the population total for the g -th combination group based on P panels. The combination groups used to construct the MCBS estimates are defined in terms of age group and initial year of Medicare eligibility (also referred to as “accretion” status). For combination group g , the a_{gp} ’s in equation (2) are generally proportional to the panel sample sizes, n_{gp} and are subject to the condition that $a_{g1} + a_{g2} + \dots + a_{gP} = 1$.

From equation (2), the variance of the estimated total can be written as:

$$\text{var}(\hat{y}) = \sum_{g=1}^G \sum_{p=1}^P a_{gp}^2 \text{var}(\hat{y}_{gp}) + D, \quad (4)$$

where D represents the total covariance between pairs of panel estimates, \hat{y}_{gp} and $\hat{y}_{g'p'}$. Although D cannot be assumed to be zero, it is expected to account for a relatively small part of the total variance. Moreover, the variance of the panel estimates can be written as:

$$\text{var}(\hat{y}_{gp}) = \frac{M^2 B_{gp}^2}{m_{NC}} + \frac{N^2 W_{gp}^2}{m\bar{n}}, \quad (5)$$

where m_{NC} = the number of noncertainty PSUs in the sample, m = the total number of PSUs in the sample, M = the number of PSUs in the population, \bar{n} = the average number of sample beneficiaries per sample PSU, and N = the number of beneficiaries in the population. B_{gp}^2 and W_{gp}^2 are unit variances associated with the different stages of selection; B_{gp}^2 is the “between PSU” unit variance and is a function of the PSU selection probabilities, while W_{gp}^2 is an average “within PSU” variance that reflects all stages of selection within the PSU (Hansen, Hurwitz, and Madow, 1953).

Although equation (5) is an oversimplification, it does serve to point out that both between-PSU and within-PSU components can change as the PSU sample ages. Since B_{gp}^2 is a function of the original PSU selection probabilities (e.g., see Hansen, et al., 1953, page 397), it can increase if the distribution of Medicare beneficiaries within PSUs changes dramatically over time. Similarly, these same changes can lead to inflated values of W_{gp}^2 due

to a redistribution of the ZIP cluster sample sizes within PSUs.

As mentioned earlier, the design effect provides a rough measure of the relative precision of the MCBS sample design with respect to a simple random sample of the same size. For example, the design effect (*DEFF*) for an estimated total is defined as

$$DEFF = \text{var}(\hat{y}) / \sigma_{SRS}^2, \quad (6)$$

where $\text{var}(\hat{y})$ is given by equation (4) and σ_{SRS}^2 is the hypothetical variance that would have been obtained from a simple random sample of the same size. In general, it is difficult to disentangle the various sources of variance contributing to the design effect. In particular, MCBS estimates are subject to both clustering and unequal weighting design effects. For estimates of means and proportions, an approximation that is useful for separating out the different effects is given by:

$$DEFF = \left\{ 1 + [cv(w_i)]^2 \right\} \left\{ 1 + (\bar{n}^* - 1)\rho \right\}, \quad (7)$$

where ρ is the intraclass correlation between beneficiaries within PSUs, $cv(w_i)$ is the coefficient of variation of the sampling weights, $\bar{n}^* = \bar{n} \left\{ 1 + [cv(n_i)]^2 \right\}$ is the average PSU sample size adjusted for varying cluster sizes, \bar{n} is the average PSU sample size, and $cv(n_i)$ is the coefficient of variation of the PSU sample sizes (United Nations, 1993). Note that formula (7) can be written as $DEFF = D_w D_c$ where $D_w = 1 + [cv(w_i)]^2$ is the unequal weighting design effect and $D_c = 1 + (\bar{n}^* - 1)\rho$ is the clustering effect.

Design effects were computed for selected characteristics derived from the 1997 Access to Care data file (see Sharma, et al., 2001). Fay’s modification of the balanced repeated replication (BRR) technique was used to compute the requisite standard errors (Judkins, 1990). Using equation (7) with $\bar{n} = 149$ and $cv(n_i) = 0.29$, the calculated design effects were then used to estimate the intraclass correlation. As shown in Table 3, the intraclass correlations range from less than 0.005 to 0.02 for the items considered. The unequal weighting design effect for the statistics in Table 3 was estimated to be $D_w = 1.22$. Thus, on average, the ratio of the clustering design effect to the unequal weighting design effect ranged from 1 to over 3.

Finally, the speculated gains in precision that could be achieved with a new PSU sample are summarized in Table 4. The design effects were calculated using equation (7) for a range of values of $cv(n_i)$. A value of $cv(n_i) = 0$ corresponds to the situation where the new PSU measure of size has controlled the PSU workloads perfectly. While this is highly unlikely in view of the panel rotation employed in

the MCBS, it does provide lower bounds on the design effects that can be achieved. On the other hand, values of $cv(n_i)$ in the range of 0.10 to 0.20 are more realistic. In this case, the reduction in *DEFFs* (as compared with the *DEFFs* in Table 3) would range from two to five percent. Although the reductions are modest, the introduction of new PSUs is expected to improve sampling precision.

Table 3. Intraclass correlations and design effects

Characteristic	ρ	<i>DEFF</i>	D_c	D_c/D_w
Poor health status	0.005	2.13	1.74	1.43
Hypertension	0.002	1.52	1.24	1.02
Difficulty bathing	0.007	2.53	2.08	1.70
Difficulty walking	0.013	3.80	3.12	2.55
Limited activity	0.003	1.87	1.54	1.26
Medicaid	0.010	3.09	2.53	2.07
Risk HMO	0.022	5.45	4.47	3.66
High school graduate	0.008	2.87	2.35	1.93
Married	0.003	1.76	1.44	1.18
Income <\$25,000	0.007	2.52	2.07	1.70

Table 4. Speculated design effects with new PSU sample

Characteristic	$cv(n_i)$			
	0.00	0.10	0.20	0.30
Poor health status	2.06	2.06	2.09	2.13
Hypertension	1.49	1.50	1.51	1.52
Difficulty bathing	2.43	2.44	2.48	2.54
Difficulty walking	3.60	3.62	3.70	3.82
Limited activity	1.82	1.83	1.85	1.88
Medicaid	2.94	2.96	3.01	3.10
Risk HMO	5.12	5.16	5.28	5.47
High school graduate	2.74	2.76	2.81	2.88
Married	1.72	1.72	1.74	1.76
Income <\$25,000	2.42	2.43	2.47	2.53

3.2 Cost and Operational Implications

There are two factors leading to increased costs associated with remaining in the original PSUs: (a) greater dispersion of the sample within PSUs, and (b) increasingly unequal workloads across PSUs. The selection of new ZIP fragments each year to represent newly-created ZIP Codes tends to disperse the sample within PSUs, thereby increasing travel costs and data collection time. The supplemental sample (i.e., new panel) selected each year also changes the relative sample sizes between the PSUs, thereby changing the individual PSU workloads. This results in additional travel and hiring costs to accommodate the changing workloads within the PSUs.

A comparison of the 1999 panel with the initial MCBS sample showed that if the 1999 panel were expanded to the size of the initial sample, the sample sizes for each PSU would fall between 89 percent and 192 percent of the original sample. Since the 1999 sample was selected using the original MCBS PSU measure of size, it provides a good indication of how the PSU workloads can fluctuate over time. Such changes would necessitate a very different staffing configuration than the original sample to

maximize efficiency. Over time, the cost of adjusting to the relative change in workloads can be absorbed into the yearly hiring and training process. The absolute costs associated with these sample disbursement changes are difficult to measure because they are offset by the overall efficiency of the data collection system. It is easy to see, however, that as new ZIP Codes are added to the sample, interviewers must travel longer distances to reach unclustered areas. Moreover, the workloads in existing ZIP fragments can become uneven. The longer the PSU sample remains in place, the more dispersed and inefficient the sample becomes. While moving to a new set of PSUs will not totally solve the dispersion problem, it will serve to lessen its impact and help maintain desired levels of operational efficiency.

4. Selection of the New PSU Sample

Based on considerations summarized in Section 3, a decision was made to reselect the sample of PSUs. In order to retain as much of the existing field operations as possible, the PSUs were selected using procedures designed to maximize overlap with the existing MCBS PSUs.

4.1 Definition of PSUs

Experience has shown that the types of PSUs defined for the MCBS and many other national in-person surveys (i.e., PSUs consisting of metropolitan areas or groups of rural counties) are generally robust and efficient for the purpose of maximizing sampling precision and minimizing survey costs. For this reason, the same PSU definitions developed for the original MCBS sample were maintained whenever possible.

In the nonMSA areas, a PSU was defined to be a single county unless it was too small to provide an adequate workload for an interviewer. In such cases, the county was combined with an adjacent county or counties to form the PSU. Each nonMSA PSU was designed to have a minimum measure of size of roughly 3,100 Medicare beneficiaries.

4.2 Certainty PSUs

For the redesign, those PSUs with at least 224,000 Medicare beneficiaries were included in the sample with certainty. (For cost reasons, Alaska and Hawaii, which together account for 0.6 percent of all Medicare beneficiaries, were not included in the sampling process.) The cutoff of 224,000 corresponds roughly to a probability of selection of 75 percent under a probability-proportionate-to-size (PPS) sample design. The use of the specified cutoff resulted in designating the 28 largest PSUs in the United States as certainties. Of these, 27 were also certainties in the original MCBS design. In addition, the largest MSA in Puerto Rico was included in the sample with certainty.

4.3 Noncertainty PSUs

The remaining (noncertainty) PSUs were grouped by Census region and MSA status (where Puerto Rico was treated as a separate “region” for sampling purposes). Within these major groups of PSUs, detailed sampling strata were formed by sorting PSUs by the percent of Medicare beneficiaries enrolled in HMO plans (and in some cases also by the percentage of minority beneficiaries), and then forming strata of roughly equal size from this sorted list. The measure of size (MOS) assigned to a PSU was the weighted sum of the number of Medicare beneficiaries in the PSU in seven age groups, where the “weights” used to calculate the MOS were proportional to the corresponding overall target sampling rates. The use of the weighted measure of size was designed to obtain self-weighting samples of beneficiaries within each of the seven age groups, while at the same time maintaining a roughly constant PSU sample size (workload) across all noncertainty PSUs (e.g., see Folsom, Potter, and Williams, 1987). Thirty-eight noncertainty strata were formed within the continental United States and one was formed in Puerto Rico. Two PSUs were then selected with probabilities proportionate to size from each stratum using procedures designed to maximize the overlap with the existing MCBS sample.

4.4 The Ernst Algorithm

To maximize overlap with the existing PSUs, the method developed by Ernst (1986) was used to select the new sample of noncertainty PSUs. In the Ernst approach, each stratum in the new design is treated as a separate linear programming (LP) problem where the objective is to maximize the unconditional overlap subject to certain constraints involving relevant selection probabilities. The results of the optimization process are then used to select the new sample. A very superficial summary of the Ernst algorithm is given below. Readers interested in the mathematical details are referred to the excellent paper by Ernst (1986).

- Step 1—All intersections (F_1, F_2, K, F_L) between a given “new” stratum and the “old” MCBS strata were identified and labeled. In each F_h , all possible old samples were listed $(s_{hi}^o, i = 1, 2, K, R_h)$ and the selection probabilities (p_{hi}) were computed for each s_{hi}^o . Similarly, within a new stratum, all possible new samples were listed $(s_j^n, j = 1, 2, K, A)$ and the corresponding selection probabilities (π_j) were computed.
- Step 2—Next, the optimal joint selection probabilities (x_{hij}) of each combination of new and old samples in

F_h were determined by maximizing the unconditional expected overlap defined by

$$\sum_{h=1}^L \sum_{i=1}^{R_h} \sum_{j=1}^A c_{hij} x_{hij}, \tag{8}$$

subject to the constraints that

$$\begin{aligned} \sum_{h=1}^L \sum_{i=1}^{R_h} x_{hij} &= \pi_j, j = 1, 2, K, A, \\ \sum_{j=1}^A x_{hij} &= p_{hi} y_h, h = 1, K, L \text{ and } i = 1, K, R_h, \\ \sum_{h=1}^L y_h &= 1, \end{aligned}$$

where c_{hij} is the conditional expected sample overlap and y_h is the probability of selecting the h -th original MCBS stratum.

- Step 3—One of the original MCBS strata was selected with probability y_h determined as part of the solution of the LP problem.
- Step 4—Finally, a new sample of PSUs was selected from the given intersection using conditional probabilities derived from the LP procedure.

4.5 Results and Workload Implications

Overall, 63 of the 107 original MCBS PSUs were retained for the new sample. The achieved overlap of 59 percent was consistent with preliminary estimates (Westat, 2000). Table 5 summarizes the results of the PSU sampling process. Also shown are estimates of the expected relative workload per PSU in the four years after the selection of the new PSU sample. The workload estimates in the table are “typical” workloads reflecting the PSU workload associated with four active panels. Due to sample attrition, the panels are not equal in size. Older panels are generally smaller in size than newer ones. The percentages shown are rough estimates intended to reflect the different sample size losses in the component panels over time.

Table 5. Distribution of MCBS PSU sample by selection status and approximate workload

Status of PSU*	No. PSUs	Relative workload (%)			
		Year 1	Year 2	Year 3	Year 4
1. C in both samples	28	100	100	100	100
2. C in new sample, NC in old sample	1	100	100	100	100
3. NC in new sample C in old sample	1	100	100	100	100
4. NC in both samples	33	100	100	100	100
5. In new sample but not in old sample	44	30	57	80	100
6. In old sample but not in new sample	44	70	43	20	0

*C = “certainty”; NC = “noncertainty”.

Under the rotating panel design, the MCBS will be operational in 151 PSUs until the beneficiary samples in the old PSUs are completely phased out of the study. For the 44 original PSUs that are not included in the new sample, there will be no new (supplemental) samples of beneficiaries. However, beneficiaries in the three most recent panels in these PSUs will continue to be interviewed for up to three years. The workload in these PSUs will start out at roughly 70 percent of the desired workload because there will be no supplemental sample to replace the oldest panel. In the ensuing two years, the workload will dwindle further to roughly 43 percent and 20 percent of the maximum workloads, respectively, as the older panels are released from the study.

At the same time, the workload in the 44 newly selected PSUs will start out at a reduced level of approximately 30 percent since they will include only the newest panel. However, with the introduction of new panels in each of the following three years, the workload will increase to 57 percent in the second year, 80 percent in the third year, and eventually to full capacity in the fourth year of operation. For the 63 PSUs that are included in both the new and original samples, the workload will be maintained at the desired 100 percent level since the annual supplement will replace the panel that is scheduled to be released under the rotating panel design.

5. Summary

The MCBS PSU sample underwent a redesign 10 years after it was first introduced in 1991 for a number of reasons. Between 1996 and 1999, design effects for the total beneficiary sample increased by an average of three percent annually. It was anticipated that if the same PSUs remained in place, further deterioration of sampling precision would occur. Continued sampling from the existing PSUs also led to unbalanced PSU workloads and a more dispersed sample within PSUs.

In order to maximize overlap with the existing MCBS PSUs, the Ernst optimization algorithm was used to select the new PSU sample. An overall 59 percent overlap between the old and new MCBS PSUs was attained. The overlap of 46 percent among the noncertainty PSUs compared favorably with the 25 percent overlap expected with independent sampling. With the redesign, the PSU measure of size was updated with current counts of Medicare beneficiaries. Stratification of the noncertainty PSUs was also enhanced through the use of relevant information on the Medicare population. The redesign will have little or no impact on existing weighting and imputation procedures. Comparing the response rates for the initial fall interview for the 2000 panel in old PSUs with those for the 2001 panel in new PSUs, the overall response rate was slightly higher in the new PSUs (87.7%) than the old PSUs (86.7%).

6. References

- Apodaca, R., Judkins, D., Lo, A., Skellan, K. (1992). Sampling from HCFA lists. *Proceedings of the Section on Survey Research Methods of the American Statistical Association*, pp. 250-255.
- Ernst, L. (1986). Maximizing the overlap between surveys when information is incomplete. *European Journal of Operational Research*, 27, pp. 192-200.
- Folsom, R., Potter, F., Williams, S. (1987). Notes on a composite size measure for self-weighting samples in multiple domains. *Proceedings of the Section on Survey Research Methods of the American Statistical Association*, pp. 792-796.
- Hansen, M., Hurwitz, W., Madow, W. (1953). *Sample Survey Methods and Theory*, New York: John Wiley & Sons.
- Judkins, D. (1990). Fay's method for variance estimation. *Journal of Official Statistics*, 3, pp. 223-240.
- Judkins, D. and Lo, A. (1993). Components of variance and nonresponse adjustment procedures for the Medicare Current Beneficiary Survey. *Proceedings of the Section on Survey Research Methods of the American Statistical Association*, pp. 820-825.
- Kish, L. (1965). *Survey Sampling*, New York: John Wiley & Sons.
- Sharma, R., Chan, S., Liu, H., Ginsberg, C. (2001). *Health and Health Care of the Medicare Population*. Data from the 1997 Medicare Current Beneficiary Survey. Rockville, MD. Westat.
- United Nations (1993). *Sampling Errors in Household Surveys*. Department for Economic and Social Information and Policy Analysis. Statistical Division, National Household Survey Capability Programme.
- U.S. Department of Labor, Bureau of Labor Statistics and U.S. Department of Commerce, Economics and Statistics Administration, U.S. Bureau of Census (2000). Technical Paper 63, Current Population Survey, *Design and Methodology*, Washington, DC.
- Westat (2000). *Medicare Current Beneficiary Survey: Evaluation of Alternative Measures of Size for Primary Sampling Unit Selection*. Technical Report prepared for Health Care Financing Administration, Rockville, MD.
- Westat (2001). *Medicare Current Beneficiary Survey: Design and Selection of the 2001 PSU Sample*. Technical Report prepared for Health Care Financing Administration, Rockville, MD.