Adaptive Matrix Sampling for the Consumer Expenditure Quarterly Interview Survey

Jeffrey M. Gonzalez
Gonzalez.Jeffrey@bls.gov*

John L. Eltinge†

Abstract
The Consumer Expenditure Quarterly Interview Survey is an ongoing panel survey of U.S. households in which sample units typically receive the same survey protocol during each interview. Because of the high burden associated with the survey request, the U.S. Bureau of Labor Statistics is exploring alternative designs that, if implemented, would change many features of the data collection process. One such alternative is adaptive matrix sampling. In general, matrix sampling involves dividing a survey into subsets of questions and then based on some probabilistic mechanism administering each to subsamples of the main sample. To potentially compensate for the resulting loss of information, as not all questions are asked of all sample units, we propose an adaptive assignment of subsampling probabilities based on data from the first interview. We use historical data to explore potential efficiency gains incurred by the use of this form of adaptive matrix sampling, develop point estimators based on simple weighting adjustments for expenditures collected under this design, and evaluate their variance properties.

Key Words: Adaptive design; Burden reduction; Multiple imputation; Sample survey; Two-phase sampling; Variance estimation

1. Introduction

1.1 The Consumer Expenditure Quarterly Interview Survey

The Consumer Expenditure Quarterly Interview Survey (CEQ) is an ongoing rotating panel survey of U.S. households in which, for each wave, all sample units are generally administered the same survey questionnaire. Each respondent is asked questions on a common set of expenditures. These expenditures are those that can be expected to be recalled for a period of three months or longer and tend to include relatively large purchases, such as for property and automobiles, and regularly occurring purchases, such as utility bills or insurance premiums. The data collected provide the basis for revising the weights and associated pricing samples for the Consumer Price Index (CPI), one of the nation's leading economic indicators, as well as a complete picture of a household's spending pattern (BLS *Handbook of Methods*, 2007).

The CEQ was designed as a personal visit interview and takes 50 to 60 minutes to complete depending on the interview. The preferred mode of data collection, at least from managerial and data quality perspectives, is personal visit; however, a substantial proportion of interviews are currently being conducted over the telephone. Safir *et. al.* (2008) point out that the percentage of cases completed by telephone has fluctuated over the years, but most recently has stabilized at about 35 percent. The increased practice of conducting CEQ interviews over the telephone has likely been made to mitigate unit nonresponse, but even so the response rate has been gradually declining over recent years. For example, response for the survey was about 80 percent in 2000, but by 2007, the annual response rate dropped to about 74 percent (BLS *Handbook of Methods*, 2007).

1.2 Motivation for Redesigning the Survey

The burden placed on respondents to accurately recall and report their entire household's expenses made in the past three months can be very high. The length of the interview poses additional concerns for respondent burden and data quality. Another concern is that survey costs are rising for numerous reasons (e.g., salary increases for interviewers, mileage, laptops, etc.). These concerns led the U.S. Bureau of Labor Statistics (BLS) to investigate alternative designs that, if implemented, would change many features of the current data collection process, such as the length of the interview and the mode of data collection. Some of these designs involve the use of matrix sampling methods to administer only subsets of questions of the current survey to sample units; thus, reducing the length of the survey and potentially addressing concerns related to burden and data quality (Gonzalez and Eltinge, 2007a; 2007b).

We recognize that modifying the current survey solely by administering a shorter questionnaire is unlikely to significantly reduce data collection costs since, in general, data collection costs are dominated by locating and con-

*U.S. Bureau of Labor Statistics, Office of Survey Methods Research, 2 Massachusetts Avenue NE, Suite 1950, Washington, DC 20212
†U.S. Bureau of Labor Statistics, Office of Survey Methods Research, 2 Massachusetts Avenue NE, Suite 1950, Washington, DC 20212

tacting sample units (Sudman, 1967). However, investigating the operational and statistical feasibility of conducting a shorter interview provides a framework for considering additional design modifications that would yield substantial cost savings. These may include, for example, switching the primary mode of data collection from personal visit to telephone.

1.3 A Potential Partial Solution - Matrix Sampling

As previously noted, matrix sampling methods involve dividing a questionnaire into subsets of questions and then administering these subsets to different subsamples of an initial sample (Gonzalez and Eltinge, 2007a; Raghunathan and Grizzle, 1995). We illustrate these methods through a relatively simple matrix sampling design shown in Figure 1. It should be noted that throughout this article, the design presented in Figure 1 will provide the general basis for illustrating the concepts of our proposal.

If we let Y represent the full vector of survey questions and S represent the initial sample, then, the rows in Figure 1, denoted by Y_k, represent specialized subsets of questions of the full questionnaire and the columns, denoted by $S_{(k)}$, represent subsamples of an initial sample. The shaded squares correspond to data that are collected, while open squares correspond to data that are not collected. For this particular matrix sampling design, each of Y and S are partitioned into $K+1$ non-overlapping subsets such that $Y = \bigcup_{k=0}^{K} Y_k$ and $S = \bigcup_{k=0}^{K} S_{(k)}$, respectively.

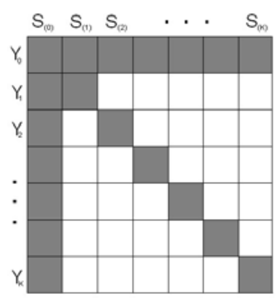

Figure 1: Illustration of a Matrix Sampling Design

In this design, we have included a "core" subset of questions, Y_0. This subset would be administered to every sample unit regardless of subsample membership (indicated by the complete shading of the first row). We have also included a subsample, denoted as $S_{(0)}$, that would receive the full questionnaire (indicated by the complete shading of the first column). Although neither feature, the "core" nor the full questionnaire subsample, is an essential component of every matrix sampling design, we incorporated the first because in the context of the CEQ, there are certain items, such as demographic characteristics (and perhaps expenditure items) that we would want to collect for all sample members. The second feature was included to accommodate public data users with various analytic objectives as well as to provide a full account of a sample unit's spending pattern.

1.4 Implementation Trade-offs

Modifications to the current sample design and questionnaire of the CEQ involve trade-offs. One example is that by shortening the length of the questionnaire, we hope to reduce respondent burden while improving data quality; however, there is a loss of information since not all questions are asked of all the respondents. To potentially account for this deficiency (and others not mentioned here), we extend simple matrix sampling methods to incorporate adaptive survey procedures. In general, adaptive survey procedures involve adjusting the survey protocol based on information collected as part of the data collection process. Therefore, we propose an adaptive form of matrix sampling in which the sample unit's subsampling probabilities are determined after some initial information is collected about the unit. The details of this proposal - including appropriate notation, estimation methods, and determination of subsampling probabilities - are presented in Section 2. The results from a simulation study in which we evaluate our method are presented in Section 3. Finally, a discussion in Section 4 concludes the article.

As a final note, for a more thorough discussion of some of the previous research on matrix sampling methods, survey design considerations, and implementation trade-offs see Gonzalez and Eltinge (2007a).

2. Adaptive Matrix Sampling

Under simple matrix sampling methods, sample units would be randomly assigned to one of the subsamples (i.e., without incorporating any information about the unit). One way to borrow from the adaptive survey procedures might be to determine these subsample assignment probabilities after some data about the sample unit has been collected. In the context of the CEQ, we propose to accomplish this by making use of the longitudinal nature of the CEQ and adjusting the probabilities of subsample assignment based on the expenditure information collected during the initial interview.

2.1 Notation

To provide the foundation for our proposed adaptive procedure of assigning subsampling probabilities, we first lay out the notation that will be used in our discussion. For $k = 0, 1 \ldots, K$, define the subsampling indicators, indicating subsample membership of the i^{th} sample unit, as $\alpha_{ik} = 1$, if $i \in S_{(k)}$. We assume that subsample assignment is based on some random process and denote the probability that the i^{th} sample unit is in the k^{th} subsample as $p_{ik} = Pr(\alpha_{ik} = 1)$. Recall from Figure 1, that the 0^{th} subsample receives the full questionnaire, so it also receives the k^{th} subset of questions. Thus, a sample unit receives subset k of the full questionnaire if and only if $\alpha_{i0} = 1$ or $\alpha_{ik} = 1$. Then the probability that the i^{th} sample unit receives the k^{th} subset of questions is the sum of p_{i0} and p_{ik}. For $k > 0$, we denote this overall probability as p^*_{ik}. Finally, let π_i be the full sample inclusion probability for the i^{th} sample unit. Define $w_i = \pi_i^{-1}$ as the inverse-probability, or design, weights and the set of modified design weights as $w^*_i = w_i/p^*_{ik}$.

2.2 Estimation Procedures

One of the population quantities of interest for the CE survey program is that of mean expenditures (e.g., the average expenditure per consumer unit on a particular item). The (full) sample estimator utilized by the survey program is the standard weighted ratio mean estimator. So, for expenditure item k, the (full) sample estimator of $\overline{Y}_k$:

$$\hat{\bar{y}}_k = \left(\sum_{i \in S} w_i\right)^{-1} \left(\sum_{i \in S} w_i y_{ik}\right). \tag{1}$$

Using the notation of the previous section, for expenditure item k, the design-based matrix sampling estimator of $\overline{Y}_k$ is given as:

$$\hat{\bar{y}}_{1k} = \left(\sum_{i \in S} w^*_i(\alpha_{i0} + \alpha_{ik})\right)^{-1} \left(\sum_{i \in S} w^*_i(\alpha_{i0} + \alpha_{ik})y_{ik}\right). \tag{2}$$

2.3 Connection with Two-Phase Sampling

For the current discussion, suppose that the first CEQ interview has been conducted but additional information on subsamples of the initial sample is desired. In our setting, this additional information will be collected approximately three months later during the second CEQ interview. This is similar to the setup in two-phase sampling. In general, for a two-phase sampling design, a sample is selected by an arbitrary sample design during the first phase and information is collected from these units. With the aid of this information, a second phase sample is selected and the key survey variables are observed for every element of the second phase sample (Särndal, et al. 1992).

The primary motivation for linking our proposed method to two-phase sampling is to explore the properties of the design-based matrix sampling estimator for the mean expenditure given above as well as identify how to adapt the subsampling probabilities for the matrix sampling design. When evaluating the moments of estimators in a two-phase sampling framework, we make explicit use of conditioning arguments. In the calculations below, E_1 and V_1 denote that expectation and variance are being computed with respect to the initial full sample selection (i.e., the phase one sample), respectively. E_2 and V_2 indicate that the expectation and variance, respectively, are being computed with respect to the subsample assignment assuming a fixed initial sample. Finally, E and V mean that the expectation and variance, respectively, are being computed over both phases of the sample selection.

Therefore, using standard linearization techniques, we have, for the expectation of expenditure item k:

$$\begin{aligned} E\left(\hat{\bar{y}}_{1k}\right) &\approx E_1\left(E_2\left(\hat{\bar{y}}_{1k} \,|S\right)\right) \\ &= E_1\left(\hat{\bar{y}}_k\right) \\ &= \overline{Y}_k. \end{aligned} \tag{3}$$

From this calculation we observe that the design-based matrix sampling estimator for expenditure item k is unbiased for both the full sample estimator and the full population mean.

Next, the variance of the mean expenditure of item k can be approximately decomposed into two components. These components reflect the variation due to the initial sample selection and the additional variation incurred due to the subsampling, respectively. Thus, we have:

$$\begin{aligned} V\left(\hat{\hat{y}}_{1k}\right) &\approx V_1\left(E_2\left(\hat{\hat{y}}_{1k}\,|S\right)\right) + E_1\left(V_2\left(\hat{\hat{y}}_{1k}\,|S\right)\right) \\ &= V_1\left(\hat{y}_k\right) + E_1\left(\left(\sum_{i\in S} w_i\right)^{-2} \sum_{i\in S} w_i^2\left(\frac{1-p_{ik}^*}{p_{ik}^*}\right)\left(y_{ik}-\hat{y}_k\right)^2\right) \\ &= N^{-2}\sum_{i\in U}\frac{1-\pi_i}{\pi_i}\left(y_{ik}-\overline{Y}_k\right)^2 + N^{-2}\sum_{i\in U} w_i\frac{1-p_{ik}^*}{p_{ik}^*}\left(y_{ik}-\overline{Y}_k\right)^2 \\ &= N^{-2}\sum_{i\in U}\left[\frac{1-\pi_i}{\pi_i} + w_i\frac{1-p_{ik}^*}{p_{ik}^*}\right]\left(y_{ik}-\overline{Y}_k\right)^2. \end{aligned} \tag{4}$$

Similarly, the covariance between mean expenditure items j and k, is given below:

$$\begin{aligned} Cov\left(\hat{\hat{y}}_{1j},\hat{\hat{y}}_{1k}\right) &\approx Cov_1\left(E_2\left(\hat{\hat{y}}_{1j}\,|S\right), E_2\left(\hat{\hat{y}}_{1k}\,|S\right)\right) + E_1\left(Cov_2\left(\hat{\hat{y}}_{1j},\hat{\hat{y}}_{1k}\,|S\right)\right) \\ &= Cov_1\left(\hat{y}_j,\hat{y}_k\right) + E_1\left(Cov_2\left(\hat{\hat{y}}_{1j}-\hat{y}_j,\hat{\hat{y}}_{1k}-\hat{y}_k\,|S\right)\right) \\ &= N^{-2}\sum_{i\in U}\frac{1-\pi_i}{\pi_i}\left(y_{ij}-\overline{Y}_j\right)\left(y_{ik}-\overline{Y}_k\right) + E_1\left(\left(\sum_{i\in S} w_i\right)^{-2}\sum_{i\in S} w_i^2\left(\frac{p_{i0}-p_{ij}^*p_{ik}^*}{p_{ij}^*p_{ik}^*}\right)\left(\hat{\hat{y}}_{1j}-\hat{y}_j\right)\left(\hat{\hat{y}}_{1k}-\hat{y}_k\right)\right) \\ &= N^{-2}\sum_{i\in U}\left[\frac{1-\pi_i}{\pi_i} + w_i\frac{p_{i0}-p_{ij}^*p_{ik}^*}{p_{ij}^*p_{ik}^*}\right]\left(y_{ij}-\overline{Y}_j\right)\left(y_{ik}-\overline{Y}_k\right). \end{aligned} \tag{5}$$

From the variance calculation given above, we see that the overall variance inflation due to the subsampling is roughly the same order of magnitude as that of the initial full sample variance. Also, the worst variance inflation will occur if $w_i\frac{1-p_{ik}^*}{p_{ik}^*}$ has a strong positive association with $\left(y_{ik}-\overline{Y}_k\right)$. Under an arbitrary matrix sampling design, we would want to reduce this second component of variation. This can be achieved by assigning expenditure item k with very high probability (thus reducing the factor, $1-p_{ik}^*$) to sample units with exceptionally large (positive) deviations, $\left(y_{ik}-\overline{Y}_k\right)$. The challenge is that this information is unavailable to us prior to the start of data collection for a single interview. However, if we assume that the expenditure information collected in the first CEQ interview is a reasonable proxy for the expenditure information that will be reported during subsequent interviews (i.e., a sample unit with a large deviation in the initial interview will likely have a large squared deviation in the second interview), then we can use the expenditure information from the first interview to base our subsampling probabilities for the second interview and perhaps, subsequent interviews. In sum, using data collected and estimates obtained from the first interview, we propose to adaptively assign sample unit's second interview subsampling probability with probability proportional to $(y_{Int1,ik}-\bar{y}_{Int1,k})$, where $y_{Int1,ik}$ is the i^{th} sample unit's expenditure value reported in interview one for item k and $\bar{y}_{Int1,k}$ is the estimated initial sample mean expenditure for item k from the first interview.

3. Simulation Study

This section describes the simulation study conducted to evaluate our proposed method of adaptively assigning subsampling probabilities for a matrix sampling design. The simulation study has two objectives. The first objective is to compare the variance attributable to initial selection of the full sample and the questionnaire section subsampling. The second objective is to compare, in terms of efficiency, the properties of the competing subsample allocation methods. To accomplish these objectives, the data already collected (from April 2007 to March 2008) using the full CEQ questionnaire were used to generate the matrix sample data set. We focused our attention on the first two consecutive interviews and only included sample units that completed these two interviews.

3.1 Simulation Setup

Although the current survey collects data on a broad set of expenditure categories, we assume a situation where the full questionnaire consists of a "core" set of questions (e.g., household demographic characteristics) and five specialized expenditure categories. These categories are clothing, insurance other than health, medical, miscellaneous, and utilities. The questionnaire is split using a matrix sampling design like the one illustrated in Figure 1. So, there are six subsamples - one which receives the entire questionnaire and five that receive the "core" and one of the five specialized subsets of questions.

We explored five allocation methods and these are summarized in Table 1 (recall that $y_{Int1,ik}$ and $\bar{y}_{Int1,k}$ are the expenditure value and estimated mean expenditure for category k from the first interview, respectively). The first was an equal allocation method. Under this method, a sample unit has an equal probability of being assigned to any one of the six subsamples (i.e., $p_{ik} = 1/6$, for $k = 0, 1, \ldots, 5$). The next method, squared deviation, required that the subsampling probabilities were proportional to the squared mean deviation of i^{th} sample unit's first interview expenditure value. To account for the possibility that some expenditures categories would naturally produce large deviations, we also computed squared relative mean deviations by dividing the i^{th} sample unit's deviation by the mean expenditure value from the first interview. Additionally, each sample unit had a one-sixth probability of being assigned the full questionnaire. The subsampling probabilities for the next two methods were constructed in a similar manner, but instead of making use of the squared deviation and the squared relative deviation we used the absolute deviation and the absolute relative deviation, respectively. One final constraint for all five methods was that we forced the subsampling probabilities for each sample unit to sum to one.

Table 1: Summary of Allocation Methods

Allocation Method	Subsampling Probabilities
Equal	$p_{ik} = 1/6, \forall k$
Squared Deviation	$p_{ik} \propto (y_{Int1,ik} - \bar{y}_{Int1,k})^2$
Squared Relative Deviation	$p_{ik} \propto \left(\frac{y_{Int1,ik} - \bar{y}_{Int1,k}}{\bar{y}_{Int1,k}}\right)^2$
Absolute Deviation	$p_{ik} \propto \lvert y_{Int1,ik} - \bar{y}_{Int1,k}\rvert$
Absolute Relative Deviation	$p_{ik} \propto \frac{\lvert y_{Int1,ik} - \bar{y}_{Int1,k}\rvert}{\bar{y}_{Int1,k}}$

For each allocation method, we carried out the simulation ($M =$) $1,000$ times. For each iteration of the five simulations we randomly assigned sample members to one of the six subsamples based on their subsampling probabilities, we then computed the mean expenditure for each of the K expenditure categories, denoted as $\hat{\bar{y}}_{m1k}$, using the design-based matrix sampling estimator of $\overline{Y}_k$ presented in equation (2). We then computed the overall simulation mean for each expenditure item mean, denoted as, $\overline{\theta}_k$, the simulation variance for each expenditure item mean, denoted as V_k simulation, and the simulation covariance, $C_{k,j}$ between expenditure means. These are given in equations (6), (7), and (8), respectively.

$$\overline{\theta}_k = \frac{1}{M}\sum_{m=1}^{M} \hat{\bar{y}}_{m1k}; \tag{6}$$

$$V_k = \frac{1}{M-1}\sum_{m=1}^{M} \left(\hat{\bar{y}}_{m1k} - \overline{\theta}_k\right)^2; \tag{7}$$

$$C_{k,j} = \frac{1}{M-1}\sum_{m=1}^{M} \left(\hat{\bar{y}}_{m1k} - \overline{\theta}_k\right)\left(\hat{\bar{y}}_{m1j} - \overline{\theta}_j\right). \tag{8}$$

As noted, one of the primary objectives is to compare the variance attributable to initial selection of the full sample and the questionnaire section subsampling. So, we estimated the variance attributable to the initial sample selection using the balanced repeated replication (BRR) method. This is the variance estimation method currently used by the Consumer Expenditure Survey Program (BLS *Handbook of Methods*, 2007).

In this method the sampled primary sampling units (PSUs) are divided into 43 strata, and the sample units within each stratum are randomly assigned to one of two half samples. To implement this method we constructed 44 replicate weights based on the half sample (HS) membership of the i^{th} unit. For each replicate α, the replicate weights are given by $w_i^{(\alpha)} = w_i\,(2\delta_{hj\alpha})$ where w_i is the original design weight of the i^{th} unit and $\delta_{hj\alpha} = 1$ if $i \in$ HS j. Then, the BRR variance of the full sample estimator of the mean of k^{th} expenditure, $\hat{\bar{y}}_k$, is given as:

$$V_{BRR}\left(\hat{\bar{y}}_k\right) = \frac{1}{A}\sum_{\alpha=1}^{A}\left(\left(\sum_{i\in S} w_i^{(\alpha)}\right)^{-1}\left(\sum_{i\in S} w_i^{(\alpha)} y_{ik}\right) - \hat{\bar{y}}_k\right)^2. \tag{9}$$

As the basis for our evaluation and comparison of the properties of the five allocation methods, we used two variance-minimizing criteria that are frequently used in optimal experimental design. They are A-optimality and D-optimality. Each of these criteria reduces the comparison of methods to a univariate functional of the simulation covariance matrix. The A-optimality criterion chooses the design, or method, with the smallest trace, or the sum of the diagonal elements, of the simulation covariance matrix. In our setting, this is equivalent to choosing the allocation method with the smallest sum of the mean expenditure variances. D-optimality is a criterion that chooses the design with the smallest determinant of the simulation covariance matrix. In other words, this criterion favors the method with the smallest generalized variance (Cornell, 1996). In addition to these two evaluation criteria, we also computed the eigenvalues of the simulation covariance matrix under each method. The eigenvalues represent the proportion of the total variation that is accounted for in each component and are useful in assessing the sensitivity of the different allocation methods.

3.2 Simulation Results

In Table 2, we present a summary of the simulation variances for each of the five allocation methods as well as the estimated Phase 1 BRR variances for each of the expenditure categories. We observe that, in general, the simulation variances for each of the allocation methods and across expenditure categories are larger than the estimation Phase 1 BRR variances. A comparison among the five allocation methods reveals that the subsampling probabilities involving the absolute deviations (both methods) seem to result in somewhat smaller variances than the remaining three methods. Also, the variances for each expenditure category under these methods are closer to the Phase 1 BRR variances.

Table 2: Simulation Variances of Mean Expenditures

Allocation Method	Clothing	Insurance	Medical	Misc.	Utilities
Equal	73.88	179.41	148.41	231.10	79.54
Squared Deviation	113.50	130.02	169.23	198.94	111.05
Squared Relative Deviation	106.27	169.88	170.18	198.93	95.46
Absolute Deviation	88.02	131.23	148.79	195.05	97.88
Absolute Relative Deviation	81.72	149.44	142.46	171.79	81.31
Estimated Phase 1 (BRR)	52.15	126.31	124.68	106.34	90.11

Table 3 offers a summary of the comparison of allocation methods based on the optimal experimental design criteria. To reiterate, with regard to the A-optimality and D-optimality design criteria, the desirability of the

Table 3: Optimal Design Criteria Based Comparison of Allocation Methods

Allocation Method	$tr(\mathbf{V})$	$det(\mathbf{V})$	λ_{max}	λ_4	λ_3	λ_2	λ_{min}
Equal	712.34	34.3×10^9	236.90	182.33	143.52	82.24	67.35
Squared Deviation	722.73	47.7×10^9	214.76	173.09	140.83	114.17	79.86
Squared Relative Deviation	740.72	52.8×10^9	208.41	180.76	164.61	108.51	78.43
Absolute Deviation	660.96	30.0×10^9	199.55	157.78	135.12	90.78	77.74
Absolute Relative Deviation	626.72	22.8×10^9	179.11	155.61	134.72	88.53	68.75
Estimated Phase 1 (BRR)	499.89	3.9×10^9	191.54	110.05	107.92	63.44	26.94

allocation method increases as the values of the trace and the determinant of the simulation covariance matrix decrease. For both criteria, the allocation method in which the subsampling probabilities are proportional to the absolute relative mean deviation appears to outperform the other four methods as it has the lowest trace and determinant values. Compared with the other methods, the values of the maximum and minimum eigenvalues also suggest that there are not large disparities in the sensitivity of assigning subsampling probabilities. At least among the designs considered, these findings suggest than an adaptive matrix sampling design in which the subsampling probabilities are proportional to absolute relative mean deviation would yield the smallest average and generalized variance.

4. Discussion

We proposed an adaptive method to determine subsample assignment probabilities for a matrix sampling design. We applied our method to the Consumer Expenditure Quarterly Interview Survey. For our application, the subsample assignment probabilities for the second interview were modified using the sample unit's expenditure values reported during the first interview. Our findings suggest that, at least among the allocations considered, an adaptive matrix sampling design in which the subsampling probabilities are proportional to absolute relative mean deviation would yield the smallest loss in efficiency for the mean expenditure estimates.

Through this research we have identified features of the proposed adaptive procedure that may warrant further development. First, recall that we only considered estimators that were based on data collected and involved a simple adjustment to the design weights. These adjustments accounted for the matrix sampling design by dividing the baseweights by the subsampling probabilities. However, we ignored the remaining steps of the weighting procedures currently employed by the CE program. These additional steps include an adjustment for subsampling in the field (this occurs when an interviewer visits a sampled address and discovers a multi-unit structure there when only one is expected), a noninterview or nonresponse adjustment, and a calibration adjustment in which the respondent sample is adjusted to match housing and population controls. The nonresponse adjustment is based on cells formed by cross-classifying the following characteristics - region of the country, household tenure (owner/renter status), consumer unit (CU) size, and race of the reference person. Weighting cells with few observations can lead to highly variable and/or large weights. When this occurs, the nonresponse adjustment cells are collapsed (following a standard set of collapsing procedures). Because a potentially smaller number of sample units will be receiving a given instrument under our proposed adaptive matrix sampling design, a re-evaluation of this weighting-cell structure and the subsequent collapsing procedures would need to occur in order to account for the smaller sample size.

In addition to making adjustments to the weighting procedures, we also plan to develop an imputation method for, at the very least, estimating mean expenditures in various categories. In our particular matrix sampling design, we identified a subsample that would be administered the full questionnaire; however, in our estimation procedures this subsample was only used in the direct estimation of each expenditure category. The potentially rich information (e.g., a full account of household characteristics and spending) of this subsample could be used in the development of imputation models to recover the information not directly collected from the remaining K subsamples. In addition, a deeper understanding of the relationship among various expenditure categories should be obtained in order to ensure that categories that are strongly correlated or predictive of each other are not assigned to the same subsample. This is to ensure that when an expenditure category is not collected from a particular subsample, the data collected on the expenditure categories from that subsample (and potentially the full questionnaire subsample) can be used to predict the missing information.

Finally, we only considered a matrix sampling design in which respondents received a common set of questions plus one other specialized subset of questions. Our motivation for assigning sample units to particular subsamples was based on the variance calculation presented in Section 2.3 (i.e., we assigned, with higher probability, sample units to receive expenditure category k if they were expected to have a large deviation from the mean). We also derived the covariance between pairs of expenditure categories in order to begin an investigation into the properties of other allocation methods in which a sample unit gets the "core" set of questions plus two (or more) specialized sets of questions. The covariance, however, involves a more complicated relationship between expenditure mean deviations and subsampling probabilities.

5. Acknowledgments

The views expressed in this paper are those of the authors and do not necessarily reflect the policies of the U.S. Bureau of Labor Statistics. The authors would like to thank MoonJung Cho, Jeanette Davis, Jennifer Edgar, Scott Fricker, Karen Goldenberg, Bill Mockovak, Adam Safir, Dave Swanson, Lucilla Tan, and Nhien To for helpful discussions of the Consumer Expenditure Interview Survey, matrix sampling and earlier versions of this paper.

REFERENCES

Bureau of Labor Statistics, U.S. Department of Labor, *Handbook of Methods*, Chapter 16, April 2007 edition, Consumer Expenditures and Income. http://www.bls.gov/opub/hom/pdf/homch16.pdf (visited August 19, 2008).

Cochran, W. (1977). *Sampling Techniques.* Third Edition, New York: John Wiley & Sons, Inc.

Cornell, J.A. (1990). *Experiments with Mixtures: Designs, Models, and the Analysis of Mixture Data.* Second Edition, New York: John Wiley & Sons, Inc.

Gonzalez, J.M. and Eltinge, J.L. (2007a). Multiple Matrix Sampling: A Review. *Proceedings of the Section on Survey Research Methods, American Statistical Association*, 3069-3075.

Gonzalez, J.M. and Eltinge, J.L. (2007b). Properties of Alternative Sample Design and Estimation Methods for the Consumer Expenditure Surveys. *Paper presented at the 2007 Research Conference of the Federal Committee on Statistical Methodology*, Arlington, VA, November 2007.

Raghunathan, T.E. and Grizzle, J.E. (1995). A Split Questionnaire Survey Design. *Journal of the American Statistical Association*, **90**, 54-63.

Safir, A. and Goldenberg, K. (2008). Telephone Effects in the Consumer Expenditure Quarterly Interview Survey. *Paper presented at the annual conference of the American Association for Public Opinion Research*, New Orleans, LA, May 2008.

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

Silvey, S.D. (1980). *Optimal Design: An Introduction to the Theory for Parameter Estimation.* London: Chapman and Hall.

Sudman, S. (1967). *Reducing the Costs of Surveys.* Chicago: Aldine.

Wolter, K.M. (1985). *Introduction to Variance Estimation.* New York: Springer-Verlag.