Potential Nonresponse Bias in the 2003 Survey of Small Business Finances

Y. Michael Yang[1], Traci L. Mach[2], Lieu N. Hazelwood[2]
National Opinion Research Center at the University of Chicago[1]
Board of Governors of the Federal Reserve[2]

Abstract

We summarize the major findings from a comprehensive nonresponse bias analysis of the 2003 Survey of Small Business Finances (SSBF). After comparing response rates by key population subgroups to assess the potential of nonresponse bias, we take four approaches to estimate the magnitude of the bias on a set of key statistics. These approaches are: (1) Comparing early responses to late responses; (2) Comparing main sample estimates and nonresponse follow-up sample estimates; (3) Comparing adjusted and unadjusted estimates; and (4) Estimating nonresponse bias analytically. We also compute relative biases and bias ratios to better measure the magnitude of the bias and especially how confidence interval coverage might be affected by the bias. We conclude by pointing out possible sample design and post-survey adjustment revisions that may benefit future surveys of similar populations.

Keywords: Nonresponse Bias, Relative Bias, Bias Ratio, Weighting Adjustments, Response Propensity.

1. Introduction

The 2003 SSBF offers a rich source of information for extensive nonresponse bias analysis. The target population of the survey included U.S. for-profit, nongovernmental, nonfinancial, and nonagricultural businesses that had fewer than 500 employees and were in operation on December 31, 2003. The sample frame was constructed from the Dun's Market Identifiers™ (DMI) file, a business database maintained by the Dun & Bradstreet Corporation (D&B). The frame was stratified into 72 strata defined by the cross of business size, census division, and urban/rural status. The stratified systematic sample initially consisted of 37,600 businesses, but only 23,798 businesses were released for screening by the end of the study.

The 2003 SSBF featured extensive nonresponse subsampling to improve response rate and reduce nonresponse bias. Nonrespondents to both screener and main interviews were subsampled for further attempts. Screener nonrespondents were subsampled at about 50 percent and main interview nonrespondents at roughly 60 percent. Both screener and main interviews were conducted over the telephone.

An analysis weight was calculated for each complete case to adjust for unequal selection probabilities and compensate for differential nonresponse. Adjustments to the base weight include sample release, eligibility, subsampling, completion rate, and so on. A weighting class adjustment procedure was used to adjust for main interview nonresponse where the weighting classes were defined by predicted response propensities. For a complete description of the weighting process, the reader is referred to Yang *et al.* (2004).

The current analysis focuses on potential bias due to nonresponse during the main interview stage. Therefore, it is based on the 6,520 businesses that completed the screener survey and were found eligible for the main interview. We chose to focus on the main interview nonresponse bias for two reasons. First, data collected from the screening interview are too limited to support substantive analysis of potential bias arising from screener nonresponse. Second, screener nonresponse appears to be mostly random based on a preliminary investigation. Noncontact is the primary source of screener nonresponse, while refusal is the dominant component of main interview nonresponse. We believe that the main source of nonresponse bias is the main interview nonresponse rather than the screener nonresponse.

Each of the 6,520 businesses involved in this analysis carries a base weight that represents the number of eligible businesses in the population. The sum of the base weights is an estimate of the size of the target population. The base weight not only reflects the original selection probability but also incorporates adjustments to screener eligibility, screener nonresponse, screener nonresponse subsampling, and screener nonresponse to the subsample. In particular, the base weight contains screener nonresponse adjustments. Therefore, to the extent that the screener nonresponse adjustments are effective in removing potential bias, nonresponse bias revealed in this analysis should be attributed to nonresponse during the main interview stage.

Section 2 describes the two nonresponse bias models that guide the analytical approaches of the later sections. Section 3 lays out the analytical approaches

to estimating nonresponse bias. Section 4 reports the results of the analyses under each approach. Section 5 discusses the analysis results, focusing on the relative magnitude of the bias and its impact on the probability of error. Finally, Section 6 concludes by suggesting possible sample design and post-survey adjustment changes that may benefit future surveys of this population.

2. Nonresponse Bias Models

Survey statisticians have proposed two statistical models that link bias in estimates with nonresponse. The basic underlying premise is that nonrespondents are different from respondents with respect to the key variables measured in the survey. The early deterministic model assumes that all members of the target population are either certain to respond or certain not to respond. Under this model, bias in the unadjusted sample mean is expressed as:

$$Bias(\bar{y}_r) = (\frac{M}{N})(\bar{Y}_r - \bar{Y}_m)$$

where the first factor on the right hand side is the proportion of nonrespondents in the population, and the second factor is the difference between the respondent and nonrespondent means (Biemer & Lyberg, 2003; Groves, 1989). Thus, the magnitude of the bias is the product of the amount of missing information and the difference in the characteristic of interest between the respondents and nonrespondents. Substantial bias will arise only when both of these quantities are sizable. If either of these quantities is small, then the bias will be small.

Empirical evidence has convinced more and more researchers that every member of population has a propensity of being a respondent. The new stochastic model therefore assumes that each member's propensity or probability to respond to a survey is a random variable that varies between 0 and 1. Under this model, nonresponse bias of the unadjusted sample mean is approximated by:

$$Bias(\bar{y}_r) \approx \frac{C(\rho, Y)}{\bar{\rho}}$$

in which

$$C(\rho, Y) = \frac{1}{N}\sum_{k=1}^{N}(\rho_k - \bar{\rho})(Y_k - \bar{Y})$$

is the population covariance between the survey variable of interest and the response probabilities (Bethlehem, 2002).

Both models hold that nonresponse bias is variable-specific; nonresponse bias depends on the amount of nonresponse; and the magnitude of nonresponse bias depends on the correlation between response propensity and the survey variable. It follows that the unadjusted estimator is unbiased if there is no correlation between the survey variable and the response propensity.

3. Analytical Approaches to Assessing Bias

Under the guidance of these two models, we employ the following analytical approaches to assess the nonresponse bias in the key estimates of the 2003 SSBF.

First, we compare the weighted (i.e., by base weight) response rates across some key population subgroups. The weighted response rate is the complement of (M/N) in the deterministic model, and it is frequently used as an indirect measure of potential nonresponse bias. Tabulating response rates by size of business is especially useful for assessing bias in a business survey. Large differences in response rates among subgroups indicate that nonresponse bias may exist.

Second, we use the level-of-effort approach by comparing two sets of respondents that required different levels of data collection effort. Specifically, we compare early respondents and late respondents on key survey variables to approximate the difference between respondents and nonrespondents $(Y_r - Y_m)$. We then multiply this difference by the weighted nonresponse rate to derive nonresponse bias estimates. This approach is effective to the extent that the late respondents are similar to the nonrespondents on the survey variables.

Third, for a more direct estimate of $(Y_r - Y_m)$, we compare the estimate from the main sample with the estimate from the nonresponse follow-up sample. Since the estimate from the follow-up sample represents the initial nonrespondents, the difference between the two estimates may be considered a closer approximation of the difference between the respondents and nonrespondents in the target population.

Fourth, we compare the nonresponse adjusted and unadjusted estimates. If the adjusted estimate is considered relatively unbiased, then this comparison provides another measure of nonresponse bias in the unadjusted estimate.

Finally, we derive the nonresponse bias analytically by estimating the two factors in the stochastic model, $C(\rho, Y)$ and $\overline{\rho}$, directly from the sample. This may be considered a novel approach because the properties of the estimates are not established. While the weighted response rate can be used to estimate $\overline{\rho}$, the estimate of the covariance term are based on the respondents only—the y variable is not available for the nonrespondents. We decided to report the results from this exercise for its heuristic value. Besides, the estimated biases under this approach are generally comparable to those under the other three approaches.

4. Results

4.1 Comparing Response Rates by Subgroups

Figure 1 shows the 2003 SSBF main interview outcome rates. The unweighted response rate is 52.4 percent and the weighted response rate is 54.4 percent. Final refusals account for about 20 percent of the sample while the "Other" category accounts for another 20 percent. Other than a small number of partial completes and cases with language barriers, the "Other" category represents nonrespondents that were not selected to the nonresponse follow-up sample and the vast majority of these cases may also be considered refusals. Therefore, the vast majority of the nonrespondents in the main interviews are refusals.

Figure 1: Main Interview Outcome Rates

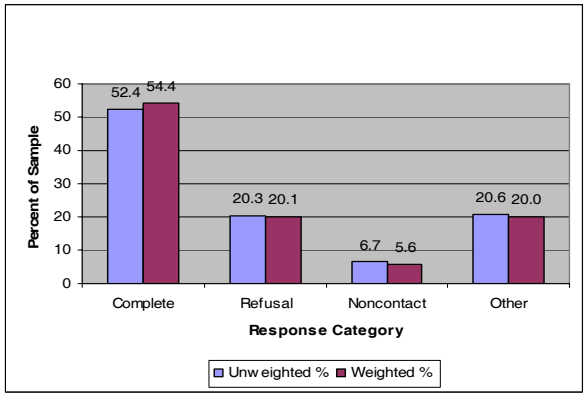

Although the target population of the SSBF includes only small businesses with less than 500 employees, the population distribution is still highly skewed in terms of business size. Nearly 94% of the businesses in the frame have less than 20 employees; and less than one percent of the businesses have more than 100 employees. Larger businesses are oversampled relative to smaller businesses to reduce errors in the estimates.

Figure 2 shows that response rate and other outcome rates vary by business size. Larger businesses are more difficult to contact and less likely to cooperate once contacted. To the extent that the survey variables are correlated with business size, this indicates a strong possibility of nonresponse bias.

Figure 2: Outcome Rates by Business Size

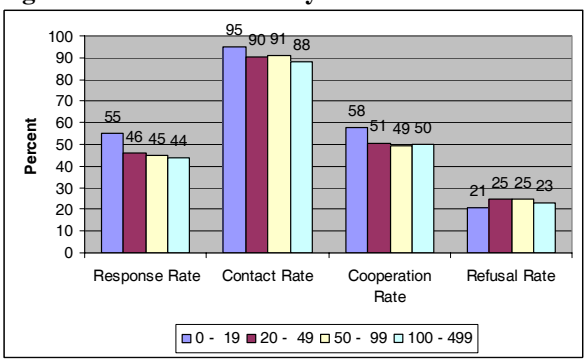

Figure 3 shows that response rates vary between urban and rural businesses. Urban businesses are slightly less likely to get contacted and much less likely to cooperate relative to rural businesses. The result is a much lower response rate for urban businesses, another indication of potential nonresponse bias.

Figure 3: Outcome Rates by Urban/Rural Location

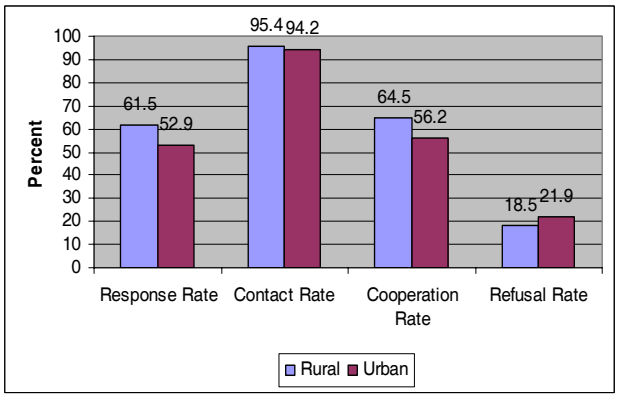

Similarly, Figures 4 and Figure 5 show substantial variation in response rates across the nine census divisions and the eight industry categories. The variation in contact rate is relatively small, so the variation in response rates is mainly due to differential cooperation rates.

Figure 4: Outcome Rates by Census Division

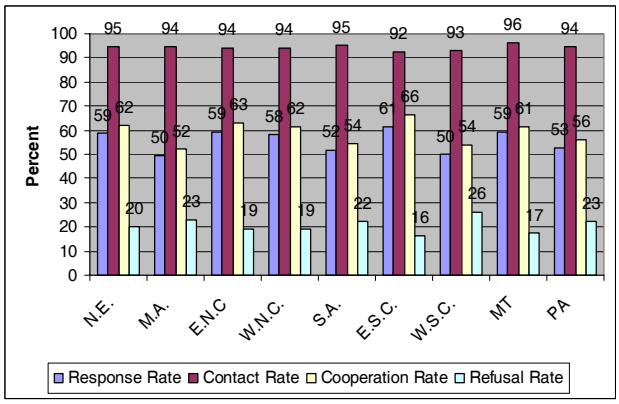

Figure 5: Outcome Rates by Industry

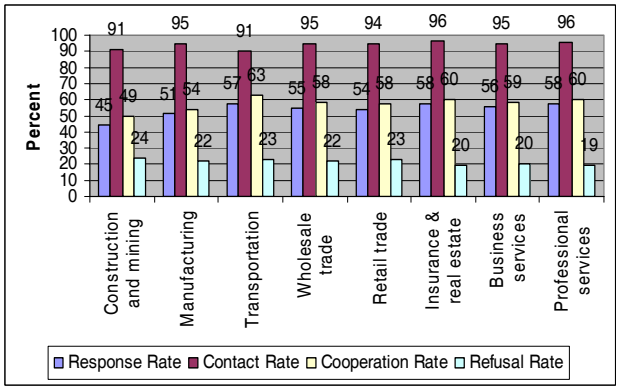

Comparing response rates across these subgroups does not yield direct estimates nor does it prove the existence of nonresponse bias on key variables. However, these comparisons reveal the differential distributions of respondents and nonrespondents on the subgroup variables. If these subgroup variables are correlated with the key survey variables, which is very likely, then the above comparisons provide a strong indication of nonresponse bias in key SSBF estimates.

4.2 Comparing Early and Late Responses

This so-called level-of-effort approach assumes that the hard to complete late respondents resemble the final nonrespondents on key survey variables. Therefore, comparing the estimates between early respondents and late respondents may be used to approximate the difference between respondents and nonrespondents, i.e., $(Y_r - Y_m)$. This approach is appealing in that it directly addresses the question of what happens to the quality of the data if interview effort were reduced to a certain level. One weakness of this approach is its lack of an objective definition of late respondents because level of effort (e.g., number of telephone calls to contact or complete a case) is usually a continuous measure. Different definitions are likely to lead to different estimates of $(Y_r - Y_m)$. The obvious guidance is the level of effort actually applied to data collection in a given survey. For the 2003 SSBF main interview, about half of the 3,287 completes were obtained with less than 10 calls, about 70 percent with less than 19 calls, and about 85 percent with less than 32 calls. In the comparisons presented below, we define those respondents that required at least 19 calls as late respondents. This definition regards the last 30% of the respondents as similar to nonrespondents.

Figures 6 and 7 compare a total of 17 key SSBF estimates between early and late respondents. All estimates are weighted by the base weight that does not contain nonresponse adjustments. The 14 estimates in Figure 6 are presented as percentages. Among them, the first 12 estimates represent the proportion of the businesses that utilize a particular financial service. For example, among early respondents, 36.4 percent of the businesses use Card Processing Services while the corresponding estimate is 38.8 percent among late respondents. The other two percentages estimate the proportion of the businesses owned by a minority or white owner. For example, 92.4 percent of the early responding businesses are owned by white while only 87.1 percent of the late responding businesses are owned by white.

Figure 6: Proportion Estimates: Early (<19 calls) vs. Late (>=19 calls) Respondents

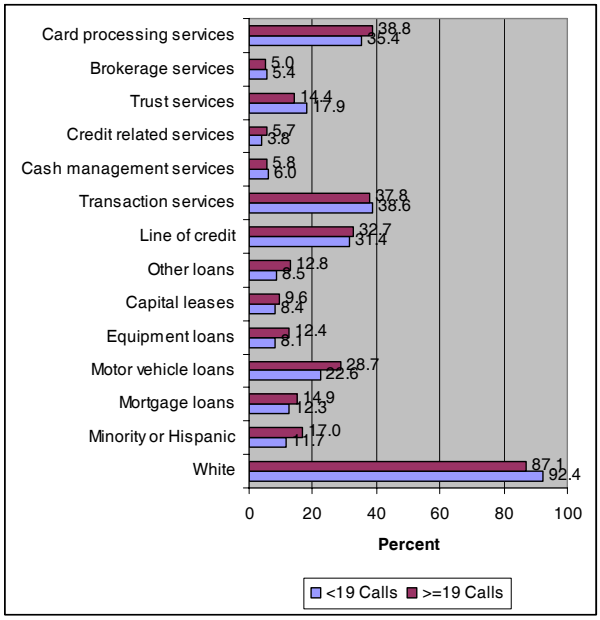

With only a few exceptions, Figure 6 shows that the estimates are significantly higher for late respondents. The only exception is associated with "Trust Services" where the observed difference is significant but in the

opposite direction. The differences in Brokerage Services, Cash Management Services, and Transaction Services are not significant.

Figure 7 presents three comparisons of estimated means between early and late respondents: Total Sales, Profit, and Total Assets. The estimated mean sales and profit are significantly higher among late respondents, while the mean assets are not significantly different between early and late respondents.

Figure 7: Mean Estimates: Early (<19 calls) vs. Late (>=19 calls)

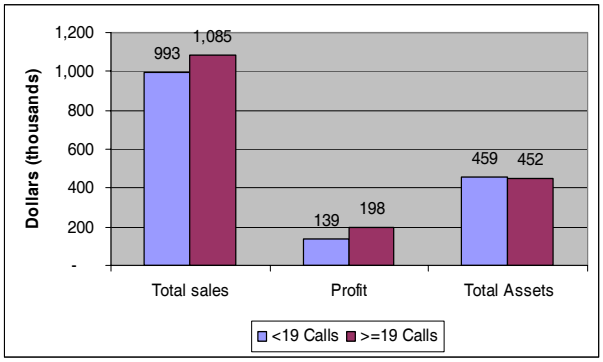

The patterns observed in Figures 6 and 7 are largely consistent with the earlier observation that larger businesses are less likely to respond to the survey than smaller ones. Late respondents are more likely to be larger businesses that tend to use more services and have greater sales, profit, and assets.

4.3 Comparing Main Sample and Follow-up Sample Estimates

Another approach to conducting a nonresponse bias analysis is to collect data from a sample of nonrespondents through intensive follow-up efforts. Estimates derived from the follow-up sample are then compared to those from the main sample. Differences between the two estimates can then be used to assess nonresponse bias. The 2003 SSBF features extensive nonresponse subsampling during the main interview stage, making it possible to conduct such direct comparisons between respondents and nonrespondents.

Figures 8 and 9 depict the comparisons between the main sample estimates and nonresponse follow-up sample estimates on the same set of 17 statistics. For the proportion estimates in Figure 8, the picture is mixed because the direction of the difference is not always consistent across the statistics. Specifically, four estimates are higher for the main sample, five estimates are higher for the follow-up sample, and the remaining comparisons are not significantly different.

Figure 8: Proportion Estimates: Main Sample vs. Follow-up Sample Estimates

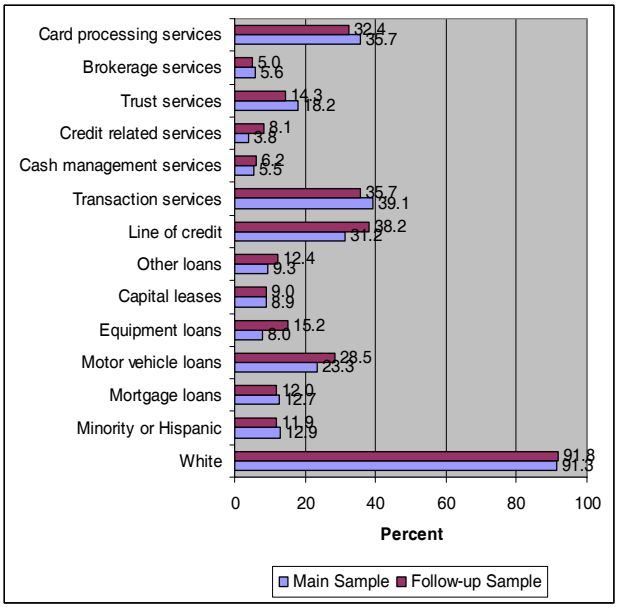

Figure 9: Mean Estimates: Main Sample vs. Follow-up Sample Respondents

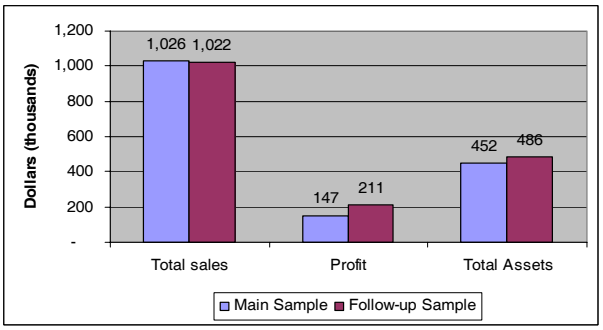

Figure 9 shows that the three mean estimates are quite close between the main sample and the follow-up sample, with only the estimated average profit showing a significant difference between the two samples.

4.4 Comparing Adjusted and Unadjusted Estimates

The comparisons in 4.2 and 4.3 are based on sample estimates that are not adjusted in any way to compensate for main interview nonresponse. In practice, the 2003 SSBF estimates are adjusted for nonresponse through weighting class adjustments where the weighting classes are defined by predicted response propensities from a logistic regression model. All the subgroup variables considered earlier, such as business size, urban/rural location, census division, industry category, and other measures available are included in the logistic model. This weighting

procedure assumes that respondents and nonrespondents in the same weighting class share similar response propensities and distributional properties on the survey variables. To the extent that this assumption is valid, the weighting class adjustments can eliminate the nonresponse bias so that the weighted estimates are approximately unbiased. Therefore, the amount of bias eliminated through nonresponse adjustment procedures may be viewed as another measure of nonresponse bias. An obvious weakness of this approach is that there is no guarantee that the adjusted estimates are unbiased or even less biased because the assumptions underlying the weighting procedures are generally not testable. For example, although cases with similar response propensities were grouped into the same weighting class, it is not knowable whether they share similar distributional properties on the key survey variables.

Figures 10 and 11 illustrate the comparisons between adjusted and unadjusted estimates on the 17 statistics. With few exceptions, nonresponse weighting led to an upward adjustment to the sample estimates although the amount of adjustment is generally small. The small adjustment could mean that the bias in the unadjusted mean is small to begin with, or it could mean that the implemented weighting adjustment procedures have limited abilities of reducing bias.

Figure 10: Proportion Estimates: Adjusted vs. Unadjusted Estimates

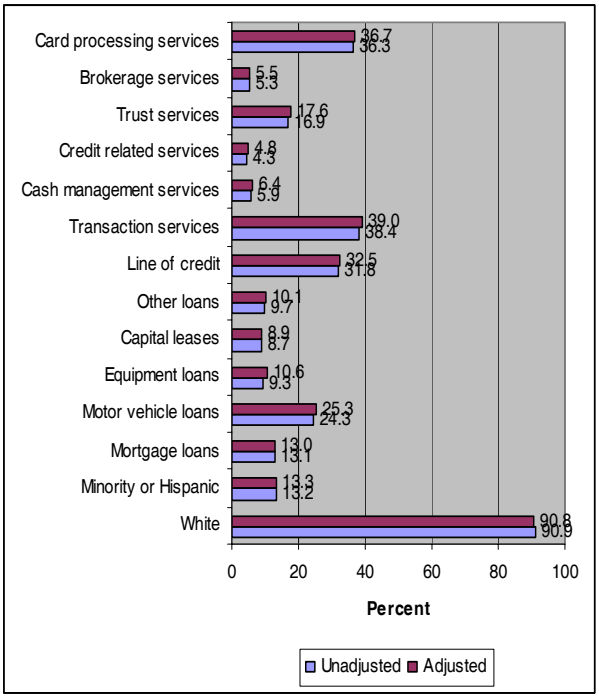

Figure 11: Mean Estimates: Adjusted vs. Unadjusted Estimates

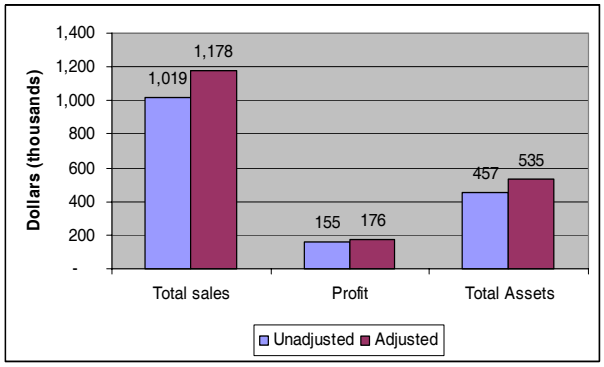

4.5 Estimating Nonresponse Bias Analytically

The stochastic model of nonresponse bias involves two population quantities: the population covariance between the survey variable and response propensity and the expected response propensity of the population. The population response propensity may be approximated by the weighted response rate. To estimate the covariance, we first fit a logistic regression model to predict a propensity score for each sample member. This model includes the following explanatory variables: number of employees, urban or rural, census division, industry category, and credit score percentile. We then computed the covariance between the 17 variables and the estimated response propensities. We finally calculated the ratio of the covariance to the mean propensity, as suggested by the stochastic model, to derive an estimate of the bias. That is, for a variable y, the estimated bias of the unadjusted mean is expressed as

$$C(\hat{\rho}, y) / \hat{\bar{\rho}}$$

The estimated covariance term is flawed because values of y are not available from nonrespondents. Nonetheless, this exercise provides a valuable heuristic approach among the alternative approaches to approximating nonresponse bias. The estimated biases are generally greater under this approach than under the other approaches considered earlier.

5. Discussion

We first compared response rates by population subgroups to assess the potential of nonresponse bias. Response rates vary significantly by business size, urban/rural location, census division, and industry category. To the extent that these characteristics are correlated with the survey variables, nonresponse could lead to biased estimates. We then took four

different approaches to estimating the magnitude of the nonresponse bias for a set of 17 key variables:

- A1—comparing estimates between early and late respondents;
- A2—comparing main sample and nonresponse follow-up sample estimates;
- A3—comparing adjusted and unadjusted estimates; and
- A4—estimating nonresponse bias analytically.

While A3 and A4 derive the estimated bias directly, the estimated differences from A1 and A2 represent $(\overline{Y}_r - \overline{Y}_m)$ in the deterministic model. Multiplying this difference by (M/N), which is estimated by $(1 - weighted\ response\ rate)$, we derived the estimated bias under A1 and A2.

Table 1 presents the bias estimates under the four approaches. The last three variables represent thousands of dollars and the rest are all percentages. Of the 68 bias estimates in the table, 55 are negative, indicating downward bias in the unadjusted mean or proportion. For 7 of the 17 statistics, the direction of the bias (negative) is the same under all four approaches. For the other 10 statistics, the direction of the bias is not always consistent although it is mostly negative. In terms of the magnitude of the bias, A1 and A2 appear to suggest greater bias than A3 and A4 for the estimated proportions; but A1 and A2 indicate smaller bias than A3 and A4 for the estimated means. Except for two small positive values, the estimated bias for average sales, profit, and assets is negative under all four approaches.

Table 1: Bias Estimates under the Four Approaches

Characteristic	A1	A2	A3	A4
Card processing services	-1.60	1.55	-0.33	-0.31
Brokerage services	0.21	0.29	-0.16	-0.11
Trust services	1.64	1.76	-0.71	-0.92
Credit related services	-0.87	-1.99	-0.53	-0.50
Cash management svcs	0.09	-0.32	-0.51	-1.11
Transaction services	0.36	1.60	-0.57	-0.80
Line of credit	-0.59	-3.26	-0.65	-1.08
Other loans	-2.00	-1.41	-0.38	-0.41
Capital leases	-0.55	-0.05	-0.18	-0.42
Equipment loans	-1.99	-3.30	-1.27	-0.54
Motor vehicle loans	-2.80	-2.38	-1.02	-0.57
Mortgage loans	-1.21	0.34	0.07	-0.22
Minority or Hispanic	-2.41	0.45	-0.10	-0.35
White	2.42	-0.23	0.07	0.16
Total Sales	-41.9	2.3	-158.6	-669.1
Profit	-27.1	-29.4	-20.8	-64.2
Total Assets	3.5	-15.6	-77.8	-31.5

A better measure of the magnitude of the bias is the relative bias, defined as the ratio of the estimated bias to the estimate. That is,

$$Relative\ Bias = B(\overline{y}_r)/\overline{y}_r$$

Table 2 shows that the relative bias is quite small for most of the statistics. The relative bias is rarely over 40 percent and mostly within 5 percent of the estimate. Many of the smallest relative biases are reported under A3, which may be an indication that the implemented weighting adjustments only corrected a portion of the nonresponse bias in the unadjusted estimate. As expected, the relative bias tends to be larger for the estimated sales, profit, and total assets due to significant correlation between business size and response propensity. The two largest relative biases are associated with Total Sales and Profit under A4.

Table 2: Relative Bias under the Four Approaches

Characteristic	A1	A2	A3	A4
Card processing services	-4%	4%	-1%	-1%
Brokerage services	4%	5%	-3%	-2%
Trust services	9%	10%	-4%	-5%
Credit related services	-18%	-41%	-11%	-10%
Cash management svcs	1%	-5%	-8%	-17%
Transaction services	1%	4%	-1%	-2%
Line of credit	-2%	-10%	-2%	-3%
Other loans	-20%	-14%	-4%	-4%
Capital leases	-6%	-1%	-2%	-5%
Equipment loans	-19%	-31%	-12%	-5%
Motor vehicle loans	-11%	-9%	-4%	-2%
Mortgage loans	-9%	3%	1%	-2%
Minority or Hispanic	-18%	3%	-1%	-3%
White	3%	0%	0%	0%
Total Sales	-4%	0%	-13%	-57%
Profit	-15%	-17%	-12%	-37%
Total Assets	1%	-3%	-15%	-6%

The presence of nonresponse bias can distort the probability of error and confidence interval coverage. The amount of the distortion depends on the bias ratio, defined as the ratio of the bias to the standard error:

$$Bias\ Ratio = B(\overline{y}_r)/\sigma_{\overline{y}_r}$$

Cochran (1977) shows that bias has little impact on confidence interval coverage if it is less than one tenth of the standard error. Even the bias ratio reaches 20 percent its impact on the probability of the total error is modest. As bias ratio increases, however, the effect becomes more serious. For example, when the bias ratio is 1, the total error rate jumps to .17, which is more than three times of the presumed .05. Furthermore, errors in the two directions are affected differently, depending on the direction of the bias.

Table 3 shows that the bias ratio is substantial for most statistics and under all approaches. Only a handful of the bias ratios are under 20 percent and many of them are over 100 percent. This means that the probability of the Type I error associated with the confidence interval of the unadjusted mean would be much greater than the nominal level. When the bias ratio is 1, bias accounts for half of the mean squared error.

Table 3: Bias Ratio under the Four Approaches

Characteristic	A1	A2	A3	A4
Card processing svcs	-138%	134%	-29%	-27%
Brokerage services	40%	56%	-31%	-21%
Trust services	189%	202%	-82%	-106%
Credit related svcs	-171%	-389%	-104%	-98%
Cash management	17%	-60%	-94%	-206%
Transaction services	31%	138%	-49%	-69%
Line of credit	-54%	-296%	-59%	-98%
Other loans	-262%	-185%	-50%	-54%
Capital leases	-79%	-7%	-26%	-60%
Equipment loans	-264%	-439%	-169%	-72%
Motor vehicle loans	-269%	-228%	-98%	-55%
Mortgage loans	-153%	43%	9%	-28%
Minority or Hispanic	-298%	56%	-12%	-43%
White	347%	-32%	10%	23%
Total Sales	-77%	4%	-292%	-1231%
Profit	-139%	-151%	-107%	-329%
Total Assets	12%	-53%	-265%	-107%

Discouraging as the sizeable bias ratios appear to be, they also reflect the fact that the estimated standard errors are very small given the large sample size. For a proportion estimate, for example, reducing the sample size by half would increase the standard error by about 41 percent, which translates to a reduction of the bias ratio by about 30 percent assuming the amount of the bias is fixed.

6. Concluding Remarks

Although the four different approaches to nonresponse bias estimation lead to different estimates of bias, the general picture that emerges from these analyses is quite clear. The target population of the 2003 SSBF is highly skewed in terms of business size. Larger businesses tend to use more services and have greater sales, profit, and assets than smaller businesses. Larger businesses also tend to have lower response rate. This significant correlation between survey variables and response propensity led to unadjusted sample estimates that are mostly downward biased. The relative bias is quite small for most statistics. However, due to the large sample size and small standard error, the nonresponse bias, if not adjusted, would have severely increased the Type I error rate.

Judging from the bias estimates under A3, the implemented nonresponse adjustments appear to be effective because the adjustments are mostly in the right direction. However, these adjustments also tend to be smaller than the bias estimated under the other approaches, which may be an indication that the nonresponse adjustments probably eliminated only a fraction of the nonresponse bias. The implication is that the weighting method may be improved in future surveys although our ability to do so is limited due to the lack of information on the nonrespondents. It is likely that variables that are not available on the frame influence both response propensity and the survey variables. In that situation, additional weighting adjustments are not possible.

One important lesson from this study is that obtaining data from large businesses is critical to the success of the survey. The 2003 SSBF oversampled large businesses. It may be advisable to apply even higher sampling rate to large businesses in the future. Furthermore, instead of selecting a random subsample of the nonrespondents for nonresponse follow-up, as was done in 2003, future surveys may consider including all nonresponding large businesses in the nonresponse follow-up sample.

Disclaimer

The views expressed herein are those of the authors. They do not necessarily reflect the opinions of the Federal Reserve Board or its staff.

References

Bethlehem, J.G. (2002). "Weighting Nonresponse Adjustments Based on Auxiliary Information," In R.M. Groves, *et al*. (Eds.), *Survey Nonresponse*, (pp. 275-288). John Wiley and Sons.

Biemer, P. P. and Lyberg, L.E. (2003). *Introduction to Survey Quality*. John Wiley and Sons.

Cochran, W.G. (1977). Sampling Techniques (3rd ed.), New York: John Wiley and Sons.

Groves, R.M. (1989). *Survey Errors and Survey Costs*. John Wiley and Sons.

Lessler, S. and Kalsbeek, W. (1992). *Nonsampling Errors in Surveys*. John Wiley and Sons.

Yang, Y.M., Rachel Harter, and Janella Chapline. (2004). "Sampling Plan for the 2003 Survey of Small Business Finances," Report submitted to the Board of Governors of the Federal Reserve, Washington DC.