RDD Telephone Surveys: Reducing Bias and Increasing Operational Efficiency

David Ferraro[1], Tom Krenzke[1], Jill Montaquila[1]
[1]Westat, 1650 Research Blvd., Rockville, MD 20850

Abstract
It is well documented that nonresponse and noncoverage rates have increased steadily in random digit dial (RDD) surveys. As these rates continue to rise, there are increasing challenges to reducing bias and increasing operational efficiency. This paper discusses several methods to consider when conducting RDD surveys. We also discuss the benefits and shortcomings of the approaches – shortcomings if it is decided not to implement a particular procedure. For example, there is a growing concern with nonresponse bias; what is the impact of by-passing the household nonresponse adjustment? We also consider the impact of cell phone only households, purging nonresidential numbers, mailings incentives, subsampling of cases prior to refusal conversion, level of effort, and other uses of auxiliary data.

Key Words: RDD, coverage, nonresponse

1. Introduction

There are many issues with Random Digit Dialing (RDD) methods that reduce efficiency and possibly induce bias in the final estimates. This includes issues such as declining response rates as well as declining coverage rates due to cell phone only households. RDD surveys are still a useful and viable means for data collection. We present several RDD approaches for putting resources where they are needed most to help reduce bias. The methods fall under two broad topics: bias due to coverage and nonresponse and operational efficiency. There is more detailed information on all of these topics in other papers (references provided), but this paper presents a brief overview of several key issues especially for new users of RDD surveys and also as reminders for ongoing users. While some of the issues may seem obvious, we have seen all of them used incorrectly or incompletely in practice.

2. Reducing Bias in RDD Surveys

2.1 Coverage of Cell Phone Only Households
The undercoverage of landline RDD samples is a growing concern to survey methodologists and practitioners. Recent estimates (Blumberg and Luke 2008) are that nationally, about 16 percent of U.S. households have only cellular service, and the rate continues to increase. Traditional landline RDD surveys, that exclude telephone exchanges designated for cell usage, do not cover this portion of the population.

The benefits of including cell phones include improved coverage especially within subgroups with higher prevalence. This improved coverage comes with a cost. It is two to three times more expensive to sample cell phone along with the additional statistical and operational considerations when implementing a cell phone sample and weighting the combined samples. Methodology is developing to include cell phones and many recent papers describe the issue in detail. In 2008, the American Association for Public Opinion Research (AAPOR) Cell Phone Task Force published the report "Guidelines and Considerations for Survey Researchers when Planning and Conducting RDD and Other Telephone Surveys in the U.S. With Respondents Reached Via Cell Phone Numbers" (available, as of October 2008, at http://www.aapor.org/uploads/Final_AAPOR_Cell_Phone_TF_report_041208.pdf). Public Opinion Quarterly has a 2007 Special Issue on "Cell Phone Numbers and Telephone Surveying in the U.S." which covers many topics related to

cell phone households and surveys. Free full text access is available at the website http://poq.oxfordjournals.org/ (current as of October, 2008).

2.2 Coverage of Landline Households

Additionally, recent research (Fahimi, Kulp, and Brick 2008) has shown that the common practice of restricting RDD sampling frames to only numbers in "1+" 100-banks (i.e., the inclusion of 100-banks, or sets of telephone numbers having a given first 8 digits, that have at least one white pages-listed telephone number; and the exclusion of 100-banks having no white pages-listed telephone numbers) results in undercoverage of almost 20 percent of households with landlines. While characteristics of cell-only households have been studied, the characteristics of households in zero-listed banks are unknown.

Issues with undercoverage of the 1+ 100-banks might be remedied by broadening the criteria for inclusion in the RDD frame. One alternative to consider is broadening the inclusion criterion to include in the landline RDD sample telephone numbers in 1,000-banks with at least one listed number. At this early stage of research, it is believed that including 1+ 1000-banks rather than 1+ 100-banks might increase the coverage rate by about by 10 percentage points, however the residential hit rate is expected to decrease by about 10 percentage points. With effective means (e.g., purging) of screening out many nonresidential telephone numbers that do not require interviewer labor, the increased cost of the lower hit rate of this approach would probably not prohibit its use. Additional research in this area is needed.

2.3 "Do Not Call" Lists

Some surveys exclude households on "Do Not Call" lists or unlisted numbers. Surveys do not have to and should not exclude these cases from their sample. With regard to the "Do Not Call" Registry, the Federal Trade Commission exempted survey research from the restrictions imposed by the Do Not Call Implementation Act. However, households on this list may not understand why they are being called. Thus, it is advisable that an appropriate response to a respondent's comment that he/she is on the "Do Not Call" list be prepared for interviewers, which may convert that respondent. The impact of the National Do Not Call Registry on response rates is largely unknown, however, recent work by Link, Mokdad, Kulp, and Hyon (2006) has shown no significant impact.

In the 2006 Washington State Population Survey (WSPS), households that were asked to be put on their "do not call" list were excluded from the survey and not called back. These cases are not necessarily on the national do not call list but are likely similar. This group was shown to be different from the eligible population in frame characteristics such as metro status, percent college graduates, tenure and income. Therefore, by excluding this group, these findings show the possibility of incurring a bias. Including either of these groups may generate a small number of complaints, which is one drawback, but that can be minimized with giving the interviewers appropriate responses.

2.4 Multiple Phone Numbers

In an RDD survey, within sampling strata or oversampling domains, each sampled telephone number has a known probability of selection. The household's probability of selection depends on the probabilities of selection of all telephone numbers in that household. If a household has more than one telephone number used for residential purposes, the household could have been sampled on any of these numbers, which has sampling and weighting implications. In order to identify the correct number, additional questions may be included in the screener to determine the number and usage of telephone numbers in each household (Roth, Montaquila, and Brick, 2001). As recommended in the paper, a weighting step should be added to adjust the weight by the number of eligible telephone numbers in a household, as determined by the survey responses to the additional questions. Based on an evaluation reported in Massey and Botman (1988), the adjustment was often capped at 2 (i.e., divide the weight by 2 if the household has two or more eligible phone numbers). But in recent years, the adjustment has been often capped at 3 (i.e., if exactly two eligible phone numbers, divide by 2; if three or more eligible phone numbers, divide by 3). A multiple telephone number adjustment accounts for the multiple chances of selecting the household and reduces the associated bias. In a recent survey, about 5.5 percent of households received weight adjustments for having multiple phone numbers.

2.5 Rostering Methods

Different rostering methods have potentially different implications for coverage rates. Westat, under contract to the U.S. Census Bureau, tested an alternative household rostering procedure for the Survey of Income and Program

Participation (SIPP). The standard procedure, which was identical to that used on most other household surveys, was to administer questions to a household respondent asking about members of the unit who "live or stay there." This approach has been criticized because terms such as "live" or "stay there" are ambiguous when applied to individuals who are either not legally related to a particular household member or who reside in more than one unit. The alternative roster questions, asking respondents to report all individuals who had stayed in the unit, for at least one night, since a specified date, was based on direct residency rules. The results showed that the alternative approach would result in better coverage by increasing the number of persons who have an ambiguous attachment to the unit.

To select one adult age 18 and older within each household for an extended interview, a method has been developed in which no questions on the household composition other than the number of adults are required in the screener. This method is is generally expected to be less intrusive than other methods that requires a complete household roster and should have less adverse effects on the response rate. Complete details on the method can be found in Rizzo, Brick and Park, 2004.

2.6 Use of Auxiliary Data in Weighting Adjustments

Auxiliary data may be used in a variety of ways in the design and selection of a sample and in estimation. First there is exchange-level data available from vendors that are derived from the Census that can be used in sampling. This enables to over-sample certain groups like minorities by selecting more exchanges that have a higher percentage of race/ethnicities that are of interest.

In an ideal survey, all the units in the inference population are eligible, found on the sampling frame and household rosters, and all those that are selected participate in the survey. In practice, neither of these conditions occur. Some eligible units have a probability of zero of being selected into the sample (undercoverage) and some of the sampled units do not respond (nonresponse). If undercoverage and nonresponse are not addressed, then the estimates from the survey may be biased. In most surveys, the weights of those who respond are adjusted to represent the undercovered units and nonrespondents. There are many sources of data that can be used in the weighting adjustments, including:

- exchange-level data that is sometimes the only data available for nonrespondents, data collected in the screener interview, and
- person and household-level data from the American Community Survey (ACS), Current Population Survey or the Decennial Census.

Nonresponse results in biases in survey estimates when the characteristics of respondents differ from those of nonrespondents. The purpose of adjusting for nonresponse is to reduce the bias. A weighting class adjustment (see Brick and Kalton, 1996) is the type of nonresponse adjustment procedure that is often used. In this procedure, nonresponse adjusted weights are computed and applied separately by cell, where a cell is defined using characteristics known for both nonrespondents and respondents. The degree of bias reduction in the adjustment is larger if either response rates or the survey characteristics are similar within the cells but differ among cells. Typically the exchange-level data, together with a few characteristics of the telephone number (e.g., whether it is listed, whether a mailable address can be attached to the telephone number), are the auxiliary data available for an RDD survey. A very powerful characteristic in predicting response propensity tends to be whether or not there is a mailable address for the phone number, especially when an advanced letter is sent. Mailable addresses and advance letters are covered in more detail later in the paper. The drawback to nonresponse adjustment is that it increases the variability of the weights and increases the sampling variance of the estimates (Kish, 1992). A nonresponse adjustment is beneficial only when the reduction in bias more than compensates for the increase in variance.

In a traditional nonresponse bias analysis, the potential for bias is evaluated using auxiliary data available for all eligible cases, and results are evaluated to see if the weighting procedures reduce the impact of nonresponse on the survey estimates. The analysis can also be used in conjunction with weighting to determine the most powerful characteristics to define the adjustment cells.

The approach to adjusting for undercoverage is somewhat different from that for nonresponse because undercovered units or persons were never given the chance to be sampled. One procedure is to use data from external sources (control totals) in a process called poststratification (Holt and Smith, 1979). This approach adjusts the surveys weights to the control total within specified adjustment classes. The primary objective of poststratification is to dampen potential

biases arising from a combination of response errors, sampling frame undercoverage, and nonresponse. A secondary objective is to reduce sampling errors.

A raking procedure is sometimes used rather than poststratification so that more auxiliary information could be employed. Raking is a commonly used estimation procedure in which estimates are controlled to marginal population totals. It can be thought of as a multidimensional poststratification procedure, because the weights are basically poststratified to one set of control totals (a dimension), then these adjusted weights are poststratified to another dimension. The procedure continues until all dimensions are adjusted. The process is then iterated so that the control totals for all the dimensions are simultaneously satisfied (at least within a specified tolerance). Brackstone and Rao (1979), Oh and Scheuren (1987), and Deville and Särndal (1992) discuss raking.

Both procedures use unit-specific variables that have known control totals which are measured in a manner consistent with the survey. The more directly related they are to the analysis unit and the key survey characteristics, the more powerful the adjustment. Person-level characteristics such as age, sex and race are desirable and often used when the analysis unit is a person. Lastly, these methods are model-based and if the model is incorrectly specified, then the adjustment will not reduce bias.

2.7 Mailings

Advance mailings are often used as a means of notifying potential participants about the study and that an interviewer will be calling. Letters should be simple, to the point and brief. It can also be helpful to provide in these letters a toll-free number and/or website for respondents, with background information about the survey. Commercial suppliers are available to match mailable addresses to samples of RDD-generated telephone numbers. Telephone numbers with mailable addresses have higher contact and cooperation rates (Brick, Judkins, Montaquila, and Morganstein, 2002). These differences are partially due to effectiveness of mailings, but also partially due to the fact that "mailable" households are different from "nonmailable" households. In the 2001 NHES, screener response rates were 75 percent for cases that were mailed to, 70 percent for mailables that were not mailed to, and 55% for nonmailables.

In addition to advance mailings, mailings can also be beneficial for refusal conversation trying to get individuals who have refused to agree to participate. This letter could be tailored to addressing the specific concerns often raised during refusal (e.g., confidentiality, the “Do Not Call” list, etc.). While there is a cost associated with all of these mailings, the increased response rate and earlier cooperation (with fewer call attempts required) might make it cost neutral.

2.8 Incentives

As response rates decline, incentives are increasingly used in RDD surveys. Two approaches have been used at the screener level. One approach is to provide the monetary incentive in the advance letter before any contact with the household. In an experiment conducted by Brick, Montaquila, Hagedorn, Roth, and Chapman (2005), initial cooperation rates increased by three percent on average when an incentive ($2 or $5) was provided with each advanced letter without adding to the overall cost of the survey. This experiment included 10 conditions with varied combinations of mailing procedures for the refusal conversion mailing (1st class and Priority Mail) and respondent incentive amounts ($0, $2, $5) for both the advance mailing and the refusal conversion mailing. The benefit of the incentive ($2 or $5) to the initial cooperation rate was 5 to 7 percentage points.

A second approach is to send the incentive only to those cases that refuse to participate. Experimental groups that included incentives were more likely to respond at first refusal conversion than those who were not sent incentives. In addition, the advance and refusal incentive treatments combined to yield higher response rates after the first refusal conversion stage for those who had received payments of $4 (two payments of $2) and those who had received $5 (in either one of the payments). Lastly, priority mail did not significantly improve initial refusal conversion rates, but did slightly improve the 2nd refusal conversion rates.

2.9 Identification of Cases for Refusal Conversion

Refusals typically comprise the majority of nonresponse in RDD surveys. A standard approach is to attempt refusal conversion on firm refusals (as well as mild refusals), excluding only hostile, abusive or threatening refusals. Telephone data collection organizations often instruct the interviewer to code the strength of the refusal. Such designations may provide information that is useful to the next interviewer who will attempt the case, or to the supervisor; however, these “strength of refusal” designations should not be used to exclude from future attempts a

particular set of refusal cases, as doing so would likely introduce bias. The designation of the strength of refusal is made by the individual interviewer and may be subjective and not consistent between interviewers. The designation can also vary by experience where less experienced interviewers tend to call cases as hard when other with more experience would call them soft. Converting refusals is a very challenging task where experienced, specially trained interviewers tend to more successful.

3. Operational Efficiency

3.1 Pre-Screening of Phone Numbers

Pre-screening involves identifying non-working and business numbers prior to the start of data collection and purging these numbers from the set to be attempted by the interviewers. Pre-screening services are available from vendors at a much lower cost per number than calling the number for the survey. As an example of this service, matches to White and Yellow Pages are used to identify nonresidential business numbers. Secondly, in a tritone test, telephone numbers are dialed to identify nonworking numbers. Any telephone number classified as a nonresidential business number or where a tritone (the distinctive three-bell sound heard when dialing a nonworking number) is encountered is considered nonworking. It is now also possible to purge cell phone numbers now when it had only been on landlines. Table 1 shows purge rates for the recent Youth Media Campaign which are typical of national surveys, though the cost, rates and methods vary by vendor. We also looked at the rates by metro status and found a slightly higher purge rate for numbers in non-metro areas. A higher purge rate will reduce the number of unproductive telephone numbers to be attempted by interviewers.

Table 1: Purge Rates by Metro Status

Metropolitan Statistical Area Status	*Business %*	*Non- working %*	*Total %*
Overall	8.5	42.5	51.0
MSA	8.9	41.2	50.1
Non-MSA	7.1	47.4	54.5

Source: Youth Media Campaign, 2006

3.2 Predictive Dialing

An alternative to purging is to use predictive dialing. Predictive dialing is performed by specialized telephone equipment that dials several telephone numbers simultaneously while a complex set of algorithms "predicts" when an interviewer will be off the phone. When used correctly, interviewers receive calls already connected to live respondents. The respondents typically hear no delay and do not know that the call was delivered to an interviewer an instant before he or she said “hello”. The labor costs are dramatically reduced because interviewers no longer receive ring no answers or busy signals.

There are drawbacks, however, to the use of predictive dialing. When used too aggressively to minimize labor costs, too many phone lines are dialed simultaneously, which results in some respondents hearing dead air after answering the telephone because no interviewer was available to take the call immediately. "Abandonment rate" is the term given to calls in which the respondent hangs up during this dead air and before the interviewer gets the call. This results in an increase in non-response. Another impact of overly aggressive predictive dialing is that the dialers will send the first answered line to an interviewer and drop the other calls. This means that respondents that can only answer after a larger number of rings have a lower probability of being interviewed. Inability to answer quickly can be correlated with health status, age, or other factors, which could introduce bias.

In addition to reducing cost, the pre-screening and predictive dialing methods may positively impact interviewer morale and save interviewing time by allowing the interviewers to concentrate primarily on more productive numbers. However, it is possible to incorrectly identify a residential number as non-working or business and have it purged from the sample. This can result in a bias but the rates are generally low.

3.3 Level of Effort

Call scheduling includes the establishment of a protocol involving proper timing of interviewing and callback procedures, including the basis for establishing the number and timing of call attempts, success rates at each level, assignments of priorities to cases based on call histories, and close and careful supervisor review. Several examples in the literature have demonstrated that people who are hard to reach are likely to be different from those who are easier to reach. In the 2001 California Health Interview Survey, interviewers made a minimum of 14 attempts to reach telephone numbers that were repeatedly busy or not answered or had answering machines. The survey industry standard is generally five to eight attempts; attempts must be on different days, different times of day and evening, and must include weekends, as well as weekdays. The chances of reaching someone at home are increased when varying the call patterns versus calling around the same time each time. It is also important to work cases equally with an appropriate number of calls. We recommend considering automation of the number of attempts per case in a call scheduler and to require a certain number of attempts based on the prior call history. Building such rules into a call scheduler might increase the number of households completed or contacted, and also help to ensure that cases are worked equally. This strategy helps put resources where they are most effective and increases the number of completed cases.

Figure 1 illustrates the importance of dialing each number an appropriate number of times. The example is from the 2006 WSPS and shows the estimated percentage Medicare received, which is a key survey outcome, by the cumulative number of calls as it converges to the final estimate. The estimates vary widely with a fewer number of calls which is a reflection of the fact that Medicare recipients may be older, home more often, and more likely to respond. The lines start to converge at 15 calls but there is still a small but statistically significant difference until 21 calls. The figure also shows that making a very large number of calls is not effective either. Although there were additional completes after 21 attempts, they had a negligible effect on the estimate and resources could have been used more efficiently aspects of the survey. It is also important to note that most characteristics did not vary by call to the extent that percent who received Medicare did, and for these other estimates, fewer than 15 calls would be sufficient.

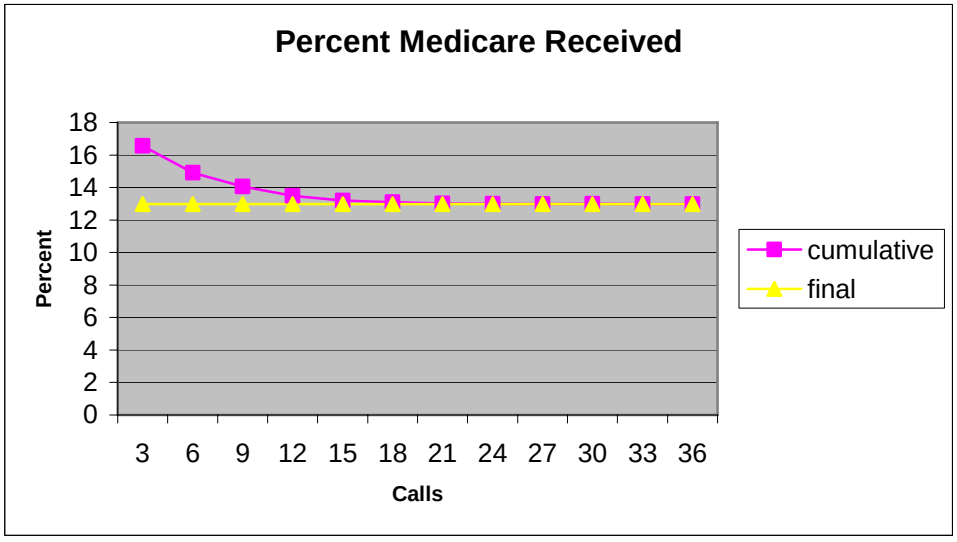

Figure 1. Estimated percentage Medicare received by number of attempts
Source: Washington State Population Survey, 2006

3.4 Subsampling Refusals Prior to Refusal Conversion

We discussed earlier the identification of cases for refusal conversion and the type of refusals to attempt. We will turn our attention to subsampling refusals so as not to attempt to convert all refusal cases. The industry standard is one conversion attempt, although in some surveys, two attempts are also made. From the standpoint of reducing bias, it is not necessary to attempt the first and/or second refusal conversion on all refusal cases, provided the cases are subsampled such that refusal cases that receive conversion attempts represent those that do not receive conversion attempts. This may be done by randomly assigning each telephone number in the sample into one of two conditions: subsampled for refusal conversion and not subsampled for refusal conversion. The subsampling rate should be arrived

at through consideration of the expected design effect introduced by subsampling; several recent surveys conducted by Westat have used a rate of 60% for the refusal conversion condition. It is beneficial operationally to assign a subsampling flag to the entire sample right after the initial selection of telephone numbers. The cases selected for the refusal attempt are for data collection first. Then the cases in the second condition, ones for which a attempt is not done, are released later. This allows extra field time for the cases where refusal conversion will be attempted.

It is important to take into account the subsampling by adding a weighting adjustment. To account for the subsampling when computing the survey weights, it is necessary to multiply the household base weight by inverse of the subsampling rate for only the cases that were retained for the refusal conversion attempt; the weights of refusal cases that were designated to not receive a refusal conversion attempt are then set to zero.

Subsampling refusals can be used to reduce the total calendar time required to conduct the survey. It has the potential to reduce data collection costs by expending less resources on those cases that take the most effort to gain cooperation and use those resources more efficiently. It may also help improve interviewer morale by reducing the number of refusals and refusal conversion attempts and this boost in interviewer morale has appeared to slightly improve response rates. The effect on estimates is minimal at best, and the weighting adjustment increases the variances slightly. The principles for refusal subsampling are well established (Hansen and Hurwitz 1946; Elliott, Little, and Lewitzky 2000) and the method has been used in surveys such as NHES, CHIS 2005 and the ACS.

3.5 Coding of Dispositions

There are many different methods used to assign final disposition, and they are not always consistent across surveys. The method used can affect a number of aspects of the survey including what numbers are redialed, residential and response rates, treatment of cases in weighting and ultimately final estimates produced from the survey. AAPOR has produced a report with a standard approach (AAPOR, 2008) for assigning dispositions. This approach not only leads to correct coding of dispositions for a particular survey but also provides consistency across surveys to allow comparisons. One especially problematic approach is basing the disposition on last call attempt. This does not account for the outcomes in intermediate calls especially if human contact is made or residency was established. Dispositions should be based on the cumulative information obtained throughout the entire sequence of call attempts.

3.6 Interactive Voice Recognition

Interactive Voice Recognition (IVR) is a technology that automates interactions with telephone callers. Advance notification via IVR is an option for nonmailables—but did not appear to have any effect when used on the 2007 NHES. A total of 14,000 nonmailable cases were designated to receive IVR call attempts; for comparison purposes 5,000 cases with address matches were included among these cases and were also mailed an advance information letter. The initial screener cooperation and refusal conversion rates for the cases in the IVR experiment were monitored on a weekly basis. With the experimental results indicating that IVR had no effect on cooperation or refusal conversion.

4. Summary

Recognizing the changing environment in the telephone research industry, research is on-going to face the major challenges of achieving acceptable coverage and response rates. While shifts in telephone survey methodology are developing, traditional RDD approaches still have a role in survey data collection. Therefore, improvements to traditional RDD approaches are needed and in this paper we have addressed a wide variety of topics and issues that arise in the design and implementation of RDD surveys. The methods implemented are often dependent on the goals and resources available to the survey, and might be affected by other aspects of the design. The choices of approaches should be based on methodological considerations, not simply on the convenience or cost of implementation.

References

Brick, J.M., and Kalton, G. (1996). Handling missing data in survey research. Statistical Methods in Medical Research, 5, 215-238.

Brick, J.M., Montaquila, J., Hagedorn, M.C., Brock Roth, S., and Chapman, C. (2005). Implications for RDD design from an incentive experiment. *Journal of Official Statistics*, 21(4), 571-589.

Brick, J.M., Judkins, D., Montaquila, J., and Morganstein, D. (2002). Two-phase list-assisted RDD sampling. *Journal of Official Statistics*, 18(2), 203-215.

Brackstone, G.J., and Rao, J.N.K. 1979. An investigation of raking ratio estimation. *Sankhya,* C (41), 97-114.

Deville, J.C., and Särndal, C.E. 1992. Calibration estimators in survey sampling. *Journal of the American Statistical Association* 87, 376-382.

Elliott, M R., Little R. J.A., and Lewitzky S. (2000). Subsampling callbacks to improve survey efficiency. *Journal of the American Statistical Association,* 95:451, 730-738.

Fahimi, M., Kulp, D., and Brick, J. M. (2008). Bias in List-Assisted 100-Series RDD Sampling. *Survey Practice,* URL surveypractice.org.

Hansen, M.H., and Hurwitz, W.N. (1946). The problem of nonresponse in sample surveys. *Journal of the American Statistical Association*, 41, 517–529.

Holt, D., and Smith, T. 1979. Poststratification. *Journal of Royal Statistical Society* A (142), 33-46.

Kish, L. (1992). Weighting for unequal Pi. *Journal of Official Statistics*, 8, 183-200.

Link, M.W., Mokdad, A.H., Kulp, D., and Hyon, A. (2006) Has The National Do Not Call Registry Helped Or Hurt State-Level Response Rates? A Time Series Analysis *Public Opinion Quarterly*, 70(5), Special Issue, 794–809.

Massey, J.T. and Botman, S.L. (1988). Weighting adjustments for random digit dialed surveys. In: Groves, R.M.; Biemer, P.P.; Lyberg, L.E.; Massey, J.T.; Nicholls II, W.L.; Waksberg, J. (eds). 1988. *Telephone Survey Methodology*. New York: John Wiley & Sons. 143-160.

Oh, H.L., and Scheuren, F.J. 1987. Modified raking ratio estimation. *Survey Methodology,* 13, 209-219.

Rizzo, L., Brick, J.M., and Park, I. (2004) A Minimally Intrusive Method for Sampling Persons in Random Digit Dial Surveys. *Public Opinion Quarterly*, 68(2), 267-274.

Roth, S.B., Montaquila, J., and Brick, J.M. (2001). Effects of telephone technologies and call screening devices on sampling, weighting and cooperation in a random digit dialing (RDD) Survey. *ASA Proceedings of the Joint Statistical Meetings.*

The American Association for Public Opinion Research. (2008). Standard Definition Final Dispositions of Case Codes and Outcome Rates for Surveys. 5th edition. Lenexa, Kansas: AAPOR.