The CHUM: A Frame Supplementation Procedure for Address-Based Sampling

Joseph McMichael, Bonnie Shook-Sa, Jamie Ridenhour, and Rachel Harter
RTI International

Proceedings of the 2013 Federal Committee on Statistical Methodology (FCSM) Research Conference

1. Introduction

Scientific surveys must be based on sampling frames that provide reasonable coverage of the target population. Address-based sampling (ABS) studies have grown in use thanks to the availability of the U.S. Postal Service's Computerized Delivery Sequence (CDS) File, which is made available to the public through non-exclusive licensing agreements with qualified vendors. The CDS file, in combination with a supplemental address list called the No-Stat file[1], contains all U.S. postal delivery points. While there is not always a one-to-one correspondence between delivery points and dwelling units, mailing addresses often serve as reasonable proxies for the physical locations of dwellings. The number of mailing addresses on the CDS and No-Stat files compares favorably with housing unit estimates from Census (Iannacchione 2011)[2]. Although these two files provide near-complete coverage of dwelling units, certain types of postal delivery points (e.g., P.O. Boxes, Rural Route Boxes) are not suitable for in-person studies because they cannot be physically located based on their mailing addresses. Addresses that are suitable for in-person fieldwork are commonly referred to as city-style addresses[3]. The combined CDS and No-Stat files allow for the central development of sampling frames for in-person surveys without the need for costly field visits prior to sample selection. ABS is increasingly viewed as a lower cost alternative to field enumeration (FE) for in-person surveys of the general population (Iannacchione 2011).

Most sampling frames have some coverage error. Traditional FE list frames contain neither the addresses of dwelling units built after the list frame was compiled nor units missed during enumeration. ABS frames for in-person surveys have been shown to suffer from undercoverage in rural areas where a higher proportion of mailing addresses are not locatable for the purposes of an in-person survey (Dohrmann et al. 2007; Iannacchione et al. 2007; O'Muircheartaigh et al. 2007). The process of allocating mailing addresses to area segments, called geocoding, can also lead to undercoverage in areas where geocoding is less accurate (Morton et al. 2007; Eckman and English 2012).

Traditional FE frames are often supplemented by a *missed dwelling unit procedure* during the interviewing stage of the study to ensure that any new dwelling units or units erroneously left off the FE list frame have a chance of selection. The most common procedure is the Half-Open Interval (HOI) method (Kish 1965) in which the interviewer searched from a sampled dwelling unit up to but not including the next dwelling unit on the ordered frame list, and any new dwelling units found in that interval are also selected into the sample. By linking the previously missed dwelling unit to the one listed just before it on the ordered list, the probabilities of selection for missed dwelling units are set to the be same as the sampled unit. This method requires a well-specified listing sequence so that the interviewer can follow the same path as the field enumerator whose information was used to create the FE frame. The problem with the HOI method for ABS is that the address lists are in mail delivery sequence order, which does not lend itself to well-defined half-open intervals because of the tendency of postal lists to cross streets and jump from block to block (McMichael et al. 2008A).

Survey researchers have developed procedures to supplement the CDS to obtain more complete address frames (Dohrmann and Sigman 2013; English et al. 2013). While some surveys are based on ABS or FE alone, most national, in-person, ABS designs now typically use a mixture of methods, depending on the expected coverage

[1] The USPS No-Stat file contains approximately seven million predominately rural city-style addresses not found on the CDS file (Shook-Sa et al, 2013).
[2] Housing unit estimates exclude group quarters units such as dormitories and group homes.
[3] City-style addresses consist of a street number, street name, city, state, ZIP code, and unit designation (as appropriate).

provided by the CDS for the selected geographies. *Enhanced listing* combines ABS with FE by using the address list for the segment and asking interviewers to update it by enumerating dwelling units missing from the CDS file prior to the selection of dwelling units as part of frame-building. Another approach is to base the sample on a *hybrid sampling frame*, which uses FE exclusively in segments where address coverage is predicted to be low and the CDS for the remainder.

RTI has developed a method that can use the CDS exclusively with a frame supplementation procedure known as *Check for Housing Units Missed* (CHUM) for improving CDS coverage (McMichael et al. 2008). Unlike enhanced listing, CHUM can be performed after sample selection and usually during data collection. It is similar to the HOI procedure in that the interviewer searches the selected dwelling unit for any missed dwelling units and then searches from the selected dwelling unit to the next dwelling unit on the frame, following a prescribed order. The interviewer also searches a subset of selected blocks to ensure that dwelling units in blocks without city-style addresses on the CDS have a chance of selection. The CHUM procedure corrects for geocoding error and gives every dwelling unit a chance of selection with a defined probability. It has the same advantage as enhanced listing in that it can be applied to geographies where the CDS provides varying levels of coverage. It has the additional advantage of being less expensive, since interviewers are searching small areas corresponding to selected dwelling units rather than entire geographic clusters. The CHUM can be implemented within an ABS-only design or in a hybrid sampling frame context.

2. The CHUM Methodology

The CHUM is designed to give each dwelling unit missing from the ABS frame a known probability of selection, theoretically providing 100 percent coverage of dwelling units. In practice, the actual frame coverage is dependent on correct implementation of the procedure by field staff. For an ABS sample, the geographic layout of dwelling units within selected area segments is typically not known at the sample selection stage. As previously discussed, ABS frames are not amenable to half-open interval supplementation procedures that require pre-specified start and end points due to the lack of geographic ordering of addresses on the sampling frame (McMichael et al 2008A). Instead, the CHUM specifies a start point and instructs field staff to follow a pre-established path-of-travel to identify dwelling units missing from the sampling frame, determining the end point when they encounter a dwelling unit that has an address on the sampling frame. Unlike the HOI, the CHUM does not rely on *a priori* knowledge of the end point of the geographic interval.

To ensure that every DU has a known probability of selection, there are two components of the CHUM: the *Check for Missed Units (CHUM1)* and the *Check for Missed Blocks (CHUM2)*. The Check for Missed Units provides coverage for dwelling units missing from the ABS frame that are located on city blocks that have ABS coverage, while the Check for Missed Blocks provides coverage for dwelling units on city blocks that have no ABS coverage. If the two components of CHUM, described below, are implemented correctly, then the ABS list frame, in combination with the CHUM, will provide coverage for all dwelling units within the selected area.

2.1. Check for Missed Units (CHUM1)

The Check for Missed Units (CHUM1) is designed to identify missed dwelling units that are located on city blocks with ABS coverage (i.e., city blocks with at least one dwelling unit whose address is included on the ABS frame). The start point for the CHUM1 procedure is the dwelling unit associated with a sampled address from the ABS frame. The CHUM has a pre-established path-of-travel which allows field staff to identify the "next dwelling unit" from the starting point, where the path-of-travel is typically a clockwise direction around the block. Field personnel, often field interviewers, are instructed not to cross streets to ensure that any missed DU can be identified from a single address on the frame. After the address of the next dwelling unit is determined, field staff checks whether or not the address is included on the ABS frame. If the address is on the ABS frame, then the CHUM1 procedure is complete—there are no missed dwelling units in the interval following the start point. However, if the address of the next dwelling unit is not contained on the ABS frame, then the address is recorded and the field staff continues to circumnavigate the block until a dwelling unit with its address on the ABS frame is located. After field staff

encounter this already listed dwelling unit (the end point), the CHUM1 procedure is complete, and missed dwelling units between the start and end points are sampled for inclusion in the study (see ***Section 2.2.3*** for further details on sampling missed dwelling units). ***Figure 1*** demonstrates the CHUM 1 procedure for a typical CHUM interval. In the example figure, if dwelling unit 1 were selected the field interviewer would travel clockwise around the block until reaching dwelling unit 2, picking up the starred dwelling unit missing from the frame. Alternatively, if dwelling unit 2 was selected for the sample, the field interviewer would pick up two missed dwelling units before reaching dwelling unit 3. If dwelling unit 3 were selected, no missed units would be detected before reaching the endpoint (dwelling unit 1).

Figure 1: CHUM Components 1 and 2

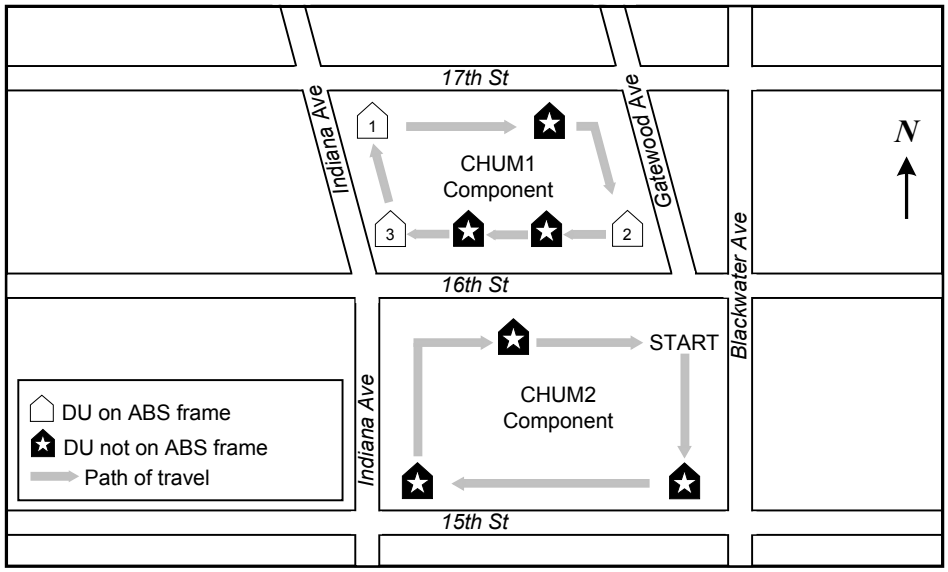

This path-of-travel works well when dwelling units are lined up along the street in a clear, clockwise geographic order (e.g., with single-family units and townhomes). However, multi-unit structures (e.g., apartment buildings) require a different path-of-travel to identify the next dwelling unit. With apartments, field staff is instructed to list the unit numbers within the building of the starting dwelling unit, to sort them alphanumerically, and to select the dwelling unit following the starting dwelling unit as the next dwelling unit. Like with the typical CHUM path-of-travel, field personnel continue listing units until they encounter one that is contained on the ABS frame (the end point).

2.2. Check for Missed Blocks (CHUM2)

If only the CHUM1 procedure was implemented, then dwelling units in polygons (usually census blocks) within the selected area with no ABS coverage would not have a chance of selection. The Check for Missed Blocks (CHUM2) procedure is designed to ameliorate this problem. During the sample selection stage, one or more census blocks within each area segment is randomly selected for the CHUM2. CHUM2 blocks can be selected based on any random sampling method. To minimize design effects, CHUM2 blocks can be selected such that the weights of missed dwelling units picked up by the CHUM2 are similar to the weights of sampled addresses from the ABS frame. An alternative approach is to select CHUM2 blocks based on probability proportional to size (PPS) sampling, giving higher probabilities of selection to blocks expected to contain large numbers of dwelling units missing from the frame. Special CHUM2 maps are created, identifying the CHUM2 start points. Start points are selected to allow coverage of all areas of the census block based on the CHUM2 path-of-travel. Multiple CHUM2 start points are sometimes required to allow complete coverage of the selected CHUM2 block. The start points are

typically selected in a clearly defined location (e.g., the intersection of two streets) so that field staff can easily determine where to begin implementing the CHUM2 procedure.

The protocol for implementing the CHUM2 procedure is the same as the protocol for implementing the CHUM1 procedure, except that the start point is a location specified on a map rather than a dwelling unit associated with an address on the ABS frame. From the start point, the field staff locates the "next dwelling unit" and checks whether or not the address of the next dwelling unit is on the ABS frame. If it is, the procedure is complete. Otherwise, the field staff continues listing and checking missed dwelling units until he/she either locate a dwelling unit with an address on the ABS frame or returns to the start point. If the CHUM2 interval contained any dwelling units associated with ABS addresses, it is not a missed area and no dwelling units are added to the frame. In these cases, any missed dwelling units in the interval are covered by the CHUM1 procedure. However, if there were no dwelling units in the CHUM2 interval that were included on the ABS frame, this was a missed area and all dwelling units in the CHUM2 interval are sampled for inclusion. ***Figure 1*** demonstrates the CHUM2 procedure. If the block at the top of the figure (bordered to the north by 17th Street and to the west by Indiana Avenue) was sampled for the CHUM2, the field interviewer would locate dwelling unit 1, which is contained on the ABS list. The CHUM2 procedure would then terminate and no dwellings would be added to the frame. However, if the block at the bottom of the figure (bordered to the north by 16th Street and to the west by Indiana Avenue) was sampled for the CHUM2, the field interviewer would circumnavigate the entire block without finding the addresses of any dwelling units on the ABS frame. This is a missed block, and all three missed dwelling units on this block would be sampled for inclusion in the study.

2.3. Sampling Issues

The probability of selection for a missed dwelling unit identified with the CHUM1 procedure is the probability of selection for the sampled dwelling unit to which it is linked. For example, if the starting dwelling unit had a probability of selection of 0.05, then all missed dwelling units identified in the CHUM1 procedure would also have a probability of selection of 0.05. However, it is not always practical to sample all missed dwelling units identified in the CHUM1 procedure. For example, if 20 missed dwelling units were identified following the sampled dwelling, subsampling could be implemented to control the overall sample size in the segment. If subsampling is implemented, the probabilities of selection must be adjusted accordingly. In the example above, if the starting dwelling unit had a probability of selection of 0.05 and 7 of the 20 missed DUs were included in the sample, the subsampled missed dwelling units would have a probability of selection of 0.05*7/20 = 0.0175.

With the CHUM2 procedure, the probability of selection for missed dwelling units is the probability of selection for the CHUM2 block. For example, if CHUM2 blocks were selected based on a simple random sample, and two of the 30 census blocks in the segment were selected, the conditional probability of selection for dwellings on missed area blocks within the selected CHUM2 block would be 2/30=0.067. (The unconditional probability is the product of this conditional probability and the probability of selecting the segment.) As with the CHUM1, dwelling units on missed blocks can be subsampled if more missed units are identified than is practical to include in the sample, and the design weights must be adjusted accordingly.

After identifying the base probabilities of selection for each sampled dwelling unit *i* (p_i), the design weight is calculated as the inverse of the probability of selection, $1/p_i$. Standard weighting procedures are then implemented on the entire sample, regardless of how the dwelling unit was identified (ABS frame vs. CHUM).

Because subsampling of missed dwelling units identified in the CHUM procedures decreases the base probabilities of selection for missed units (and therefore increases their design weights), too much differential subsampling across

the clusters can increase the variation in weights which can increase the design effect due to weighting[4] and thereby lower the precision of certain estimates. This must be considered when implementing the CHUM in areas with very low ABS coverage because subsampling will likely be required to control the overall sample size in these areas. For example, if a sampled segment has only 5 percent ABS coverage, then we would expect 95 percent of the dwelling units to come into the sample via the CHUM. This can lead to challenges in controlling the sample size and sampling inefficiencies due to unequal weighting, so care should be taken when designing the study to account for the effects of missed dwelling units on the precision of survey estimates.

3. CHUM Operational Issues

While the CHUM methodology theoretically provides complete coverage for dwelling units missing from ABS sampling frames, it is dependent on correct implementation of the procedure by field staff. The specific implementation of the CHUM can be tailored to the needs and limitations of each study. All of the following aspects of the CHUM can be optimized based on the goals and budget of the study: the selection of CHUM1 addresses and CHUM2 starting points, when the CHUM is implemented, CHUM materials provided to field staff, field staff training and support, processing missed dwelling units identified with the CHUM, and quality control procedures. The following subsections discuss operational issues associated with implementing the CHUM. These subsections include case studies of how the CHUM has been implemented previously, as well as alternative implementations for future studies.

3.1. Selecting CHUM1 Addresses and CHUM2 Points

The CHUM does not need to be completed for all sampled addresses and/or segments in the study. If a study does not have the resources available to implement the CHUM everywhere, or if the target population resides in areas where address undercoverage is less of a concern (e.g., urban areas), then the CHUM1 procedure can be completed for a subset of sampled addresses. However, implementing the CHUM1 for only a subset of addresses can impact the unequal weighting effect, as any added addresses will have a different weight than the original sample due to the subsampling of CHUM1 addresses. The more subsampling that occurs, the higher the variation in design weights between the selected addresses and addresses identified with the CHUM unless subsampling can be managed across clusters in a way that reduces or eliminates the impact on weight variation. If resources are limited, the CHUM can be focused on areas where CDS undercoverage is expected. CHUM1 can be completed for a higher proportion of addresses in segments that are expected to have undercoverage. CHUM2 blocks can be selected in a similar fashion. Depending on the resources available, the expected CDS coverage, and the sample design, any number of CHUM2 blocks can be selected for inclusion.

In addition to tailoring the number of CHUM1 and CHUM2 intervals, researchers can tailor when the CHUM procedure is implemented. The CHUM can be implemented anytime following sample selection of addresses. Field staff can be instructed to make a trip to sampled segments prior to screening and interviewing to complete the CHUM or the CHUM can be implemented during data collection. Implementing the CHUM prior to data collection allows time to develop the final sample of addresses before the start of screening and interviewing. This can allow researchers to determine appropriate subsampling rates and simplifies subsequent visits to the segment by field staff. However, additional resources are required to visit segments prior to data collection, and the results are less timely.

[4] The design effect due to weighting is equal to one plus the relative variance of the sample weights, and is a measure of how unequal weighting impacts the efficiency of survey estimates (Kish, 1965; Valliant et al, 2013).

3.2. Creating CHUM Maps

Area probability surveys require maps to help field interviewers navigate to and between selected addresses. With ABS studies, the first task an interviewer completes in a segment may be to determine whether the selected addresses are actually within the segment boundaries.

The materials field interviewers are provided for completing the CHUM are flexible. In general, they need maps that identify segment boundaries and the location of CHUM2 start points. Specialized maps can be created for the CHUM2 points showing only the blocks the interviewer needs to circumnavigate. Depending on the resources available for the study, these maps can be paper or electronic.

To implement the CHUM, field interviewers also need a method for searching the ABS frame of the sample segment for potential missed dwellings. This methodology can be as simple as a printout of addresses in that segment or as complex as an application that allows interviewers to search for addresses electronically. Electronic applications can be programmed to pre-populate candidate next addresses based on geographic proximity to the field interviewer's location or the CHUM2 starting point. If this methodology is employed, care should be taken to avoid confirmation bias by FIs – that is, quality control measures should be employed to ensure that the FIs are not treating the list as correct and not updating it (Eckman and Kreuter 2011).

3.3. CHUM Training

RTI studies that have used the CHUM to date have had trainings that lasted anywhere from thirty minutes to a full day, depending on the resources available and the expected reliance on the CHUM. A short training has lower up-front costs but will likely require more field support and more procedures for verifying quality of the frame augmentation process. Longer training is more expensive up front, but FIs should be better-prepared for what they encounter in the field. FIs can be trained on many of the rare situations they may encounter or they can be taught the basic CHUM approach while stressing to seek field support when they encounter something unusual.

3.4. Support for Field Interviewers

While performing the CHUM, field interviewers will encounter situations that require them to seek help. If their supervisor is not able to determine the correct course of action, the case may get routed back to the home office for support from someone who understands the theoretical underpinnings of the CHUM. Field interviewers sometime require help determining if the path-of-travel is taking them out of the segment; the FI may also need help to determine if the address is associated with a missed dwelling in a CHUM interval (e.g. if the address of the dwelling they encounter is similar but not identical to an address on the ABS frame). When field interviewers encounter a large number of missing units , a sampling statistician will need to establish the subsampling rate and associated procedures in that interval to mitigate the cost of all the additional cases. It is also important to have quality control checks in place; if an FI lists more units in a segment than are expected based on the predicted ABS coverage in that area, for example, perhaps he or she is off-course or not implementing the procedure correctly. How many units "more" is defined by collection resources and areas visited in a particular study – it may be acceptable for field staff to list up to three before seeking additional help or it may be acceptable for staff to list up to twenty.

3.5. Quality Checks for CHUM Implementation

Prior research has shown that interviewers do not always complete field work correctly, whether implementing the HOI, an enhanced listing procedure, or the CHUM (Eckman and Kreuter 2011; Eckman and O'Muircheartaigh 2011; Iannacchione et al. 2012). One method for monitoring the quality of CHUM fieldwork is to create situations where the field interviewer should find at least one missed dwelling unit by removing addresses from the field interviewer's address list that are likely to be the next dwelling unit. While this *seeding* method is not perfect, it allows sampling staff to have a measure of how well the CHUM is being implemented. Seeding does have cost implications, however, as it puts an additional burden both on interviewers and sampling staff to CHUM and verify unnecessary intervals. The data quality benefits and validation that the CHUM is being implemented correctly typically outweigh these costs.

More sophisticated monitoring techniques are possible when the field interviewers have Global Positioning System (GPS) capabilities in the field. Field interviewers' locations can be monitored to ensure they are following the correct path-of-travel. Mistakes can be corrected quickly enough that the procedure can be repeated in areas where it was not implemented correctly. GPS can also ensure that missed dwellings identified through the CHUM are contained within the sampled segment.

4. Advantages and Disadvantages of ABS with CHUM

As discussed in the introduction, there are other options for reducing undercoverage in ABS studies. Traditional FE and using the CDS alone with no supplementation procedures are approaches used in somewhat unique situations. Studies that use the CDS alone would either have to be restricted to areas where address coverage is quite high (typically urban areas) or would have to be willing to accept undercoverage in some (typically rural) areas. Conducting FE for all segments in a large study is both expensive and time-consuming. A study using the hybrid FE and ABS design may not have a strong need for frame supplementation in segments where the predicted coverage rate[5] is high.

Studies that do want to supplement the coverage of ABS frames in some manner have used either Enhanced Listing or the CHUM. Timeline, segment size, data collection budget, and impact on variances are all considerations in choosing the frame supplementation approach. Enhanced listing is completed prior to sample selection while CHUM is completed in the field at the start of screening and interviewing. Studies with an aggressive data collection timeline may not be able to add the extra time enhanced listing requires. Enhanced listing may be a better option in studies that are largely in areas where a large number of addresses are likely to be added to the frame. Including them ahead of time allows the statistician more control over the probabilities of selection associated with the added addresses. This can lead to more efficient sample designs.
While enhanced listing requires field staff to canvas the entire segment to supplement the CDS file, the CHUM is only implemented from randomly selected starting points. Because the CHUM is therefore less time consuming for field staff, it may be the better choice for studies with large segment sizes.

Some ABS studies use census block groups (CBGs) as the segments of dwelling units. CBGs average around 500 households – checking all addresses associated with each CBG in the sample could be quite time-consuming relative to completing the CHUM1 interval for only the sampled addresses within that CBG.

[5] The coverage rate is the predicted household coverage for a particular segment. The coverage threshold is the minimum rate at which the ABS frame is used in place of FE.

Both timeline and segment size are inherently related to budget constraints. Surveys typically require smaller segments when data collection costs are tight. Extra time spent listing the entire segment in a separate trip prior to screening and interviewing is also an extra cost.

Both supplementation procedures are dependent on field staff to implement the procedure correctly. Enhanced listing has been shown to suffer from confirmation bias (Eckman and Kreuter 2011). Without proper interviewer training, the CHUM can also lead to undercoverage (Iannacchione et al. 2012). Regardless of the method selected, resources must be available to appropriately train, monitor, and support field staff,

5. Conclusions

ABS is increasingly the preferred approach for in-person area probability studies. Although the combined CDS and No-Stat frames cover all Postal delivery points, they do not currently provide full housing unit coverage for in-person surveys. Several approaches have been developed to supplement the ABS frame where needed including hybrid sampling frames, enhanced listing, and the CHUM. Enhanced listing provides more statistical control for sampling added addresses, as researchers can select the sample of addresses from a complete frame that has already been supplemented, eliminating the need for subsampling that often occurs with the CHUM. However, the CHUM provides some operational advantages over enhanced listing. The CHUM occurs at the start of screening and interviewing rather than at the frame-building stage, so address updates are timelier. This also avoids multiple trips to the segment which saves resources. The CHUM is only implemented from a subsample of addresses in each segment, which makes it less time-consuming than enhanced listing, where field interviewers canvass the entire segment looking for missed addresses. In addition, the CHUM allows for geographic larger segment sizes compared to methods that canvas the entire segments. This has the likely benefit of lowering the intraclass correlation and thus decreasing design effects and variance (Valliant et al. 2013). Regardless of which supplementation technique is implemented, it is critical for field staff to be appropriately trained, monitored, and supported which performing field work. While the logistics of these procedures can be tailored to the needs and resources of the study, they are critical to maintaining high coverage and ensuring data quality.

References

Dohrmann, S., D. Han and L. Mohadjer (2007). Improving coverage of residential address lists in multistage area samples. *Proceedings of the American Statistical Association, Section on Survey Research Methods*.

Dohrmann, S. and R. Sigman (2013). Using an Area Linkage Method to Improve the Coverage of ABS Frames for In-Person Household Surveys. *Federal Committee on Statistical Methodology Research Conference*, Washington, D.C.

Eckman, S. and N. English (2012). "Creating Housing Unit Frames from Address Databases Geocoding Precision and Net Coverage Rates." Field Methods **24**(4): 399-408.

Eckman, S. and F. Kreuter (2011). "Confirmation bias in housing unit listing." Public Opinion Quarterly **75**(1): 139-150.

Eckman, S. and C. O'Muircheartaigh (2011). "Performance of the Half-Open Interval Missed Housing Unit Procedure." Survey Research Methods **5**(3): 125-131.

English, E., K. Dekker and C. O'Muircheartaigh (2013). Correcting Coverage Deficiencies in Address-Based Frames: The Use of Enhanced Listing. *Federal Committee on Statistical Methodology Research Conference*, Washington, D.C.

Iannacchione, V., K. Morton, J. McMichael, D. Cunningham, J. Cajka and J. Chromy (2007). Comparing the coverage of a household sampling frame based on mailing addresses to a frame based on field enumeration. *Proceedings of the American Statistical Association, Section on Survey Research Methods*.

Iannacchione, V. G., J. P. McMichael, B. E. Shook-Sa and K. B. Morton (2012). A Proposed Hybrid Sampling Frame for the National Survey on Drug Use and Health - Final Report. Research Triangle Park, NC, Report prepared for Substance Abuse and Mental Health Services Administration by RTI International.

Kish, L. (1965). Survey Sampling. New York, Wiley & Sons.

McMichael, J. P., J. L. Ridenhour and B. E. Shook-Sa (2008). A robust procedure to supplement the coverage of address-based sampling frames for household surveys. *Proceedings of the American Statistical Association, Section on Survey Research Methods*.

Morton, K., J. McMichael, J. Cajka, R. Curry, V. Iannacchione and D. Cunningham (2007). Linking mailing addresses to a household sampling frame based on census geography. *JSM Proceedings*.

O'Muircheartaigh, C., E. English and S. Eckman (2007). Predicting the relative quality of alternative sampling frames. *Proceedings of the American Statistical Association, Section on Survey Research Methods*.

Valliant, R., J. A. Dever and F. Kreuter (2013). Practical Tools for Designing and Weighting Survey Samples. New York, Springer.