Using Digital Imagery to Update the Measures of Size of Area Segments

Sylvia Dohrmann[1], Lee Harding[2], Lin Li[1]
[1]Westat, 1650 Research Boulevard, Rockville, Maryland 20850
[2] Macro International, 11785 Beltsville Drive, Suite 300, Calverton, Maryland 20705

Abstract
In area household surveys, a multi-stage sampling approach is frequently used to create a nationally representative sample. The first stage is the selection of primary sampling units (PSUs) consisting of counties or groups of counties. The second stage of selection involves the formation and selection of smaller units within the PSUs, often referred to as segments, consisting of blocks or groups of blocks. In the United States, segments formed for area samples are typically based on population or housing counts for the specific blocks or groups of blocks, as collected in the most recent decennial census.

Late in the decade, demographic or housing data from the last decennial census are likely to be inaccurate in areas with considerable growth or demographic shifts since the census taking. While intercensal updates are available at the county level, and the advent of the American Community Survey (ACS) provides updated data at the block-group level for some areas at varying intervals during the decade, population and housing unit information at the smaller block level is available only once every decade. If data used to calculate the segment measure of size (MOS) are inaccurate, the sampled segments may not yield the expected number of sampled units.

This paper describes the effect of an inaccurate MOS on the ultimate sample and the resultant accuracy of the results, along with two approaches that have been used to compensate for this inaccuracy: building permit sampling and two-phase sampling. The latter method requires that segments be selected in two phases and the MOS updated with housing unit counts obtained in the field. This paper also discusses three methods of obtaining updated housing unit counts for the two-phase approach: the original method of in-field, or windshield, canvassing; the option of replacing in-field canvassing with a purely digital method; and a hybrid method. The three canvassing methods are then evaluated based on their effectiveness and cost.

Key Words: Area sample, MOS, probability proportional to size sampling

1. Area Segment Sampling Approaches Late in a Decade

Samples selected with probability proportionate to an inaccurate MOS may yield a sample of unpredictable size and workload. Inaccurate segment MOS have intrinsic and/or extrinsic costs. Approaches that have been used to compensate for this inaccuracy include building permit sampling (Bell, et al., 1999) and two-phase sampling (Montaquila et al., 1999, 2002). Both methods have their advantages.

With building permit sampling, a separate segment frame of building permits issued since the most recent decennial census is created, in addition to the usual frame of segments formed with an MOS based on data from the last decennial census. This method has been proven effective for reducing the variation in segment sizes which, in turn, leads to improved fieldwork predictability, cost efficiency, and precision. It is not without its disadvantages, however. Not all areas can benefit from this method since they may not require the issuance of building permits for new construction. In other areas, obtaining permit information may be difficult as it will likely require that field staff visit the permit office; offices with permit information stored only in hard-copy form would further require that staff transcribe permit information by hand. Also, permit issuance is not affirmation that construction has occurred, and is not required for units such as mobile homes. In addition, since newer construction would be sampled via the permits, newly constructed units (those constructed since the last decennial census) must be screened out in regular area segments; this screening is sometimes problematic since it is often the case that residents do not know the age of their homes.

The two-phase approach introduced by Montaquila et al. in 1999 and evaluated in 2002, requires that a larger sample of segments be selected in PSUs estimated to have significant growth since the decennial census. Field staff are sent to the segments to canvass the area and estimate the number of dwelling units (DUs) in the segments. The MOS is then updated to reflect this estimate, and the final sample of segments is selected from the first phase sample based on the updated MOS.

The evaluation performed in 2002 concluded that the two-phase method improved the coverage of persons residing in newly constructed DUs. The authors also concluded that with this approach, ideally, a very large number of first phase segments should be selected, and that their number should be proportionate to the segment growth variability in the PSU. In practice, growth variability is extremely difficult to determine a priori. Unit growth is not uniform across a county, or even across smaller geographic areas such as cities or Census designated places. Additionally, the size of the first phase sample is limited operationally by the cost of sending in-field canvassers.

Beginning in late 2005 we used the two-phase segment sampling approach, with some modifications from the method documented in Mohadjer et al., 2002 for a national area household survey. (The modifications are detailed in the Appendix.) A major change from the original implementation of this approach was to keep the number of segments constant across PSUs in light of scheduling and cost concerns. However, in some PSUs, this limitation resulted in a single segment dominating the final sample due to the growth found during canvassing. The impact of any single segment could be lessened with a larger first phase sample. This limitation was one reason for our investigation of the effectiveness of in-field canvassing.

2. Methods for Updating Segment MOS

The original method of in-field, or windshield, canvassing requires that field staff drive through the segments and count the number of DUs within the segment boundaries. This is meant to be a rough count, and not an exact enumeration; a full enumeration is conducted at the listing stage (when the addresses of all units in the final sample of segments are captured). While windshield canvassing is less time consuming than the listing process, it requires time and travel.

The introduction of geospatial and satellite digital imagery applications allows users to view images of an area without travel costs. In many instances, it is possible to discern individual DUs, and compare the number of units in an area to the decennial census count, or, in essence, “digitally canvass” the area.

In the following sections we review the performance of the original windshield canvassing method; consider the performance of a purely digital canvassing method; and present the results of a hybrid approach which utilizes the advantages of both methods.

2.1 Windshield Canvassing

To assess the performance of the windshield canvassing method, we compared the number of DUs found by the canvassers with that found through the more rigorous listing process. The number captured by the listers is our best estimate of the actual number of units in the segments. Hence, the performance of windshield canvassing can only be evaluated for the segments sampled at the second phase as only these were listed. The windshield canvassing method was used in 18 PSUs containing a total of 432 second phase segments.

One advantage of canvassing segments is that growth and decline in an area may be captured and incorporated into the MOS. By contrast, building permit sampling allows only for the estimate of growth. To determine whether the windshield canvassing was effective in capturing change and, thus, improving the accuracy of the segment MOS, we compared the difference between the census and lister counts to the difference between the canvass and lister counts, to see which was closer to the listers’ more accurate count:

$$\frac{\left(\left|lister - census\right| - \left|lister - canvas\right|\right)}{lister} \qquad (1)$$

For this evaluation, we consider a cut-off of 10 percent. That is, if the measure indicated by (1) above is less than or equal to 10 percent then we conclude that canvassing did little to improve the segment MOS. If the measure is greater than 10 percent, we conclude that canvassing found some change from the census and improved the segment MOS.

Table 1: Evaluation of Windshield Canvassing

		Segments with improved MOS	
	Total segments	*Number*	*Percent*
All segments	432	159	36.8
Segments with at least 50% change	83	77	92.8

The results of our review of the windshield canvassed segments are shown in Table 1. While windshield canvassing improved the MOS of 37 percent of the segments, in the remaining 63 percent of the segments, windshield canvassing did not improve the MOS. Thus, in most of the segments, canvassing was not necessary. The larger concern, however, is the performance of windshield canvassing in segments that experienced a large amount of change. Among the 432 segments, there were only 83 segments that experienced at least a 50 percent change in either direction. In virtually all of these segments, windshield canvassing accurately captured that change (77 segments or 93%).

Since the two-phase approach is only implemented in PSUs known to have significantly changed since the decennial census, we cannot conclude that the results in Table 1 indicate that it was necessary to canvass in only 37 percent of the *PSUs*. Rather, the segments were distributed *across* the PSUs. Given the cost of sending field staff to perform windshield canvassing, it was disappointing to find that this procedure was not necessary in most of the segments evaluated. This led us to investigate other means of updating segments' MOS.

2.2 Digital Canvassing

Over recent years, the quality of aerial and satellite images available to the public have improved drastically. The resolution of some images allows users to see details such as houses, sidewalks, and cars. These images are available from the United States Geological Survey, or through applications such as TerraServer and Google Earth.

Google Earth Pro is a licensed software application which displays satellite images of varying resolution of the Earth's surface, and also allows the import of geographic information system data onto the images. This additional capability allows us to overlay segment boundaries onto the images and visualize the area within the segment. The degree of resolution available with this product is based somewhat on the points of interest, but most land is covered in at least 15 meters of resolution.

Figure 1. Satellite image with segment boundary overlay.

Figure 1 shows an image from Google Earth Pro for a sampled segment. With an image of the quality of Figure 1, the number of DUs can be determined easily by counting the rooftops (digitally canvassing the segment). In this example, the number of housing units from the decennial census exactly matches the digital count and that found by the listers.

We tested this method of digital canvassing in five PSUs containing a total of 150 second phase segments using the measure indicated by (1) above. The results are shown in Table 2. Using this measure, we found that digital canvassing improved the MOS in only 28 percent of the segments compared with 37 percent using windshield canvassing. Canvassing in this manner was considerably faster (approxmately 20 segments were canvassed per day) and much less expensive (approximately 15% the cost of windshield canvassing). The performance in segments with at least 50 percent change was markedly worse than with windshield canvassing. Segment change was effectively captured in 71 percent of these segments, compared with 93 percent captured with windshield canvassing.

Table 2: Evaluation of Digital Canvassing

		Segments with improved MOS	
	Total segments	*Number*	*Percent*
All segments	150	42	28.0
Segments with at least 50% change	34	24	70.6

The reasons for the discrepant performance between the two methods result from some inherent shortcomings of the digital images. First, the age of the image is not always apparent. While they are generally no more than three years old, if change occurred after image creation, it cannot be captured through the digital canvassing process. In some cases, areas of new construction may be apparent, indicating that the number of DUs has likely increased since image creation. Second, not all images have the required resolution to discern houses from other buildings, or houses from each other. Third, the number of units in apartment buildings or other multiple-DU structures cannot be determined from the images.

2.3 Hybrid

Given that both windshield and digital canvassing have their disadvantages, we developed a hybrid method which incorporates the best of both approaches. This method entails digital canvassing in all the first phase segments, and windshield canvassing in areas where a digital count cannot be determined (due to resolution of the image, the presence of apartment buildings, or apparent construction). The steps of the process are as follows:

- Attempt to digitally canvass all segments;
- If a full digital count does not appear possible, then update the MOS with a windshield canvass count;
- If a full digital count is possible, and it does not substantially differ from that obtained from the decennial census, then do not update the MOS; and
- If a full digital count is possible, is substantially different from the decennial census figure, and appears reasonable, then update the MOS with the digital canvass count.

This hybrid method was used in the same five PSUs and 150 segments for which we tested digital canvassing alone. Using the steps outlined above, we were able to reduce the number of segments which needed windshield canvassing by more than half. The digital images for 80 of the 150 segments were such that the count from the digital canvassing was clear and reasonable, so that visiting the segment was not necessary. For the remaining 70 segments, an in-person visit was necessary, as the digital image either was not clear, showed evidence of construction in progress, or contained multiple-unit structures.

Table 3 contains the results of the hybrid method in terms of the measure indicated in (1). The hybrid approach performs much better than digital canvassing alone, with 37 percent of the segments having improved MOS, as opposed to the 28 percent found with digital canvassing. This figure is comparable with what was found with windshield canvassing alone.

Table 3: Evaluation of Hybrid Canvassing Method

		Segments with improved MOS	
	Total segments	*Number*	*Percent*
All segments	150	56	37.3
Segments with at least 50% change	34	30	88.2

By digitally canvassing all the segments initially we were able to identify some segments which showed very little or no change in size since the decennial census, so they did not need their MOS updated. Not all of the segments given to windshield canvassers experienced large amounts of change; in fact, only 26 of the 70 windshield canvassed segments experienced more than 50 percent change. However, the extent of the change could not be captured by digital canvassing. In all but one of the few digitally canvassed segments that had more than 50 percent change from the decennial census, the windshield count accurately reflected this change.

2.4 Other Methods Considered

When investigating more efficient ways to update the MOS for segments, methods other than digital canvassing were researched, but eventually eliminated. These included using residential address lists originating from the United States Postal Service (USPS) and incorporating image recognition software into the process. While both of these possibilities were quite intriguing, they were each found to be disadvantageous when it came to updating the MOS at the segment level.

Pattern recognition technology has been used to develop software that attempts to allow a computer to “see” an image similar to the way a human might. Software using this technology may identify particular objects in images, or match two images of the same object. For use in two-phase sampling, image recognition software was investigated as a means of identifying individual DUs and then determining the total number of DUs within segment boundaries. In order for this software to be able to identify a DU, it must be given instructions on how a DU may appear in the image. This may be a simple task for a particular housing development, or even an entire segment, as the rooftops of DUs are likely to be similar in these areas. However, retraining would be necessary when moving into a new area. Separating commercial structures from residential ones are additional complications, as are multi-unit structures. With these considerations in mind, it was determined that having a human digitally canvass the segments and provide contextual comments about commercial and multi-unit structures was a better use of resources.

We also experimented with using residential address lists from the USPS in place of windshield canvassing. Since 5-digt ZIP codes are generally the smallest geographic unit for which these lists may be purchased, we first purchased all addresses in the ZIP codes covering the first phase segments. It was then necessary to obtain the geo-coordinates of each purchased address in order to determine which addresses fell into the selected segments. However, it is not always possible to geocode all addresses. As a result the updated DU counts were the summation of the number of cases geocoded into the segment and an estimated fraction of non-geocodable cases.

The biggest advantage of using purchased lists in place of windshield canvassers lies in the potential savings in cost, despite the cost of the address lists themselves and the amount of home-office time needed to create the DU estimates based on the lists. Also, the purchased lists may contain addresses that have been recently demolished. The disadvantages include undercoverage of DUs receiving mail via PO boxes only, DUs along rural routes, and group quarters. Geocoding is another complication, particularly in rural and/or fast growth areas. Since the geocoding process requires matching a particular address to a geographic information system database, an address may not be geocodable if it is created after the database was established. Hence, in faster growing areas there are likely to be a larger percentage of nongeocodable cases, making their allocation to segments even more challenging.

Previous research has indicated that in urban areas with slow development, using address lists in place of enumerators may reduce survey costs. (See Dohrmann, et al., 2007 for more information regarding the use of USPS address lists in area sampling.) However, our evaluation showed that in fast growth urban areas, there is no guarantee that address lists are updated promptly enough to reflect the most recent new construction or demolition of DUs. We concluded that the characteristics of a PSU which cause two-phase sampling to be necessary, particularly its fast growth, make using address lists for updating the MOS of first stage segments in two- phase sampling inefficient.

3. Cost Savings

Since digitally canvassing segments can be performed in the home office, this method is much less expensive than windshield canvassing. For the PSUs we examined, digital canvassing was 85 percent less than the cost of windshield canvassing. As noted above, however, the performance of this method makes it less desirable despite the dramatic cost savings.

In the PSUs we investigated, the hybrid method was slightly over half as costly as windshield canvassing (55%). The substantial cost savings of this method is a function of the number of segments that can be resolved digitally. In an area for which the images are not clear, have numerous segments with multi-unit structures, or evidence of new development in progress, the cost savings will be lessened. However, it is unlikely that the hybrid method would ever result in a higher cost than windshield canvassing alone. The only instance in which this might happen is if none of the first phase segments can be reconciled digitally. In that case, the cost of digitally canvassing would be in addition to windshield canvassing. Again, this scenario is unlikely.

4. Summary and Conclusions

Towards the end of the decade, samples which rely on decennial census data to determine measures of size (MOS) for probability samples risk basing the selection of units on inaccurate data. The result of this inaccuracy at the segment level may be the number of ultimate sampling units differing from expectation, increased variation in the number of sampling units across segments, and/or large variations in the probabilities of selection affecting the precision of the estimates. Building permit sampling and two-phase segment sampling are two methods that have been documented as ways to compensate for this inaccuracy. In this paper, we focussed on the two-phase approach.

Two-phase segment sampling requires that a larger sample of segments be selected at the first phase. The MOS for the first phase segments sampled is then updated based on more recent estimates of the number of dwelling units (DUs) in the area. The method of obtaining this updated value was originally planned to be windshield canvassing, but the limitation of the number of segments that can be efficiently canvassed caused us to investigate this method's effectiveness.

Our investigation led to the finding that only 37 percent of the first phase segments experienced enough change to warrant an update to the MOS. Hence, canvassers visit over 60 percent of the segments needlessly. Digital canvassing, viewing satellite and aerial photographs of the segments, and then counting the number of discernable DUs in the images was considered, but found to be less effective than windshield canvassing. Segments that were well established, with no multi-unit structures, and for which the images were clear could be accurately canvassed digitally. In others, digital canvassing provided an indication of change in the segment, but an accurate count could not be made simply from the image.

This observation led us to the development of the hybrid method. After digitally canvassing all the segments, only those segments for which the number of units in the segment was not discernable from the image were given to windshield canvassers. The hybrid method performed just as well as windshield canvassing alone, and was 55 percent of the cost. This cost savings will allow us to increase the first phase sample size, especially in those PSUs estimated to have grown considerably since the decennial census.

Appendix. Changes in the Implementation of the Two-Phase Approach Since 2002

The two-phase approach as originally presented by Montaquila, et al. in 1999 was implemented only in areas for which it was known, based on the most current county-level estimates available from the U.S. Census Residential Building Permits Survey, that the PSU had grown more than 10 percent since the decennial Census. Using building permit data at the Census place level from this same survey, estimates of growth were made in each Census designated place in the PSU (for those places having permit-issuing offices); these were used as estimates of the segment-level population change before selecting the first phase segment sample.

Let U_p denote the number of units for which building permits were issued since the most recent census in place *p*. The number of persons residing in newly constructed units in place *p* may be estimated by $T_p^{[u]} = 2.6U_p$. The estimate of the place-level "growth" ratio is then given by

$$g_p = \frac{T_p^{[o]} + T_p^{[u]}}{T_p^{[o]}},$$

where $T_p^{[o]}$ is the total population for place *p* from the most recent decennial census. For the first phase of segment selection, the MOS for each segment is adjusted for that segment's place-level growth. That is, the MOS for segment *i* in the first phase segment selection is M_i', where

$$M_i' = M_i g_p, i \in p.$$

Since segments are much smaller than places, estimated place-level growth may differ considerably from the true change in the size of the segment since the last decennial census. In order to obtain more accurate estimates of the true change in segment size, updated counts of the number of DUs in each of the first phase segments are obtained through windshield canvassing. These updated counts are then used to compute MOS for the second phase segment selection such that the overall probabilities of selection of the segments are accurate.

Let U_i' denote the number of DUs found by windshield canvassing first phase segment *i*. The change in the size of first phase segment *i* is estimated based on U_i' as follows:

$$g_i' = \frac{U_i'}{U_i^{[0]}},$$

where $U_i^{[0]}$ is the number of DUs in segment *i* at the time of the most recent decennial census.

The MOS for the second phase of selection is

$$\begin{aligned} M_i^{[2]} &= \frac{M_i g_i'}{M_i'} \\ &= \frac{g_i'}{g_p}, \quad i \in p \end{aligned}$$

or the ratio of the segment growth to place growth.

The evaluation of this method in 2002 (Mohadjer, et al., 2002) found that place growth estimates from the Building Permits Survey overestimated the actual growth in the first phase segments. Places are much larger than segments and new housing developments or multi-unit structures will result in areas of concentrated growth. Thus the estimate of growth in a place is not evenly distributed across all segments within that place. The impact of this is that if the growth is overestimated, the first phase MOS for a segment is also overestimated, and its probability of selection is too large. The resulting segment, if selected, will not be as large as expected.

As a result of this evaluation, rather than adjusting the first phase MOS, M_i, by the estimated place growth, g_p, the segments were simply made larger to allow for the potential growth by inflating the minimum MOS by 1.25. This value was developed based on empirical data for several PSUs so that the variation in growth may be absorbed by the final sample. The MOS for the second phase of selection is now $M_i^{[2]} = g_i'$.

The two-phase method was further modified, for operational reasons, to select a constant number of segments in the first phase across PSUs rather than varying the number based on estimated PSU growth. This was set initially to twice the number of second phase segments, and then increased to four times the number of second phase segments.

References

Bell, B., Mohadjer, L., Montaquila, J., and Rizzo, L. (1999). Creating a frame of newly constructed units for household surveys. *Proceedings of the Survey Research Methods Section*. Washington, DC: American Statistical Association, 306-310.

Dohrmann, S., Han, D., and Mohadjer, L. (2007). Improving Coverage of Residential Address Lists in Multistage Area Samples. *Proceedings of the Survey Research Methods Section of the American Statistical Association*, 3219-3226.

Montaquila, J.M., Bell, B., Mohadjer, L., and Rizzo, L. (1999). A methodology for sampling households late in a decade. *Proceedings of the Survey Research Methods Section.* Washington, DC: American Statistical Association, 311-315.

Mohadjer, L., Montaquila, J.M., Sherris, E. (2002). Evaluation of a Two-phase Approach to Segment Selection for Area Probability Samples Late in the Decade. *Proceedings of the Survey Research Methods Section.* Washington, DC: American Statistical Association, 2395-2399.