# Predicting the Relative Quality of Alternative Sampling Frames

**Colm O'Muircheartaigh**, Edward English, Stephanie Eckman
Statistics and Methodology, NORC, 55 E. Monroe St., Suite 2000, Chicago, IL, 60603

**Abstract**

The paper presents a systematic review of the relative efficacy of traditional listing and the USPS address list as sampling frames for national probability samples of households. NORC and ISR collaborated to compare these two national area-probability sampling frames for household surveys. We conducted this comparison in an ongoing survey operation which combines the current wave of the HRS with the first wave of NSHAP. Since 2000, survey samplers have been exploring the potential of the USPS address lists to serve as a sampling frame for probability samples from the general population. We report the relative coverage properties of the two frames, as well as predictors of the coverage and performance of the USPS frame. The research provides insight into the coverage and cost/benefit trade-offs that researchers can expect from traditionally listed frames and USPS address databases.

KEY WORDS: USPS lists, sampling frames, Tailored samples

## 1. Introduction

Previous work at NORC, beginning in 2001, on the evaluation of sampling frames for probability samples of households has led to a reassessment of field listing as the "gold standard" in terms of coverage and accuracy (O'Muircheartaigh, Eckman, and Weiss, 2002). NORC has been carrying out an examination of the alternative approach of using the United States Postal Service (USPS) list as a basis for frame construction for area probability surveys. In 2004, NORC and ISR embarked on a national benchmark comparison, whose goal was to provide a quantitative analysis describing the benefits and drawbacks of traditional listing vs. the USPS list for a national household sample.

In our first report on this research, we compared a traditionally-listed housing unit (HU) frame to a USPS-based frame in the same set of areas (O'Muircheartaigh, Eckman, English, Lepkowski, and Heeringa, 2005). When discrepancies arose between the two frames, however, it was not possible to determine the source of the error, or which frame was more accurate.

Since then we have conducted additional field work in a subset set of areas to reconcile the two frames with the reality on the ground. The "best" frame produced by this field effort provided a basis for determining the performance of the two frame-construction approaches, USPS and traditional listing, and allowed us to distinguish the circumstances under which each approach may be preferable. Our second report (O'Muircheartaigh, English, Eckman, Upchurch, Garcia, and Lepkowski 2006) compared the "best" frame to the traditional and USPS-derived lists. We concluded that the USPS-derived list was a better representation of reality than the traditional list in most cases. One feature of our analysis was that a priori expectations as to which frame would be superior were frequently not correct.

Our current research takes the analysis a step further by examining how well the best frame is captured by the T and U address frames in different real-world circumstances. These findings enable us to predict the success of T and U based on a-priori data, such as urbanicity, Census measures, etc. Our hope is that this research will help survey designers choose the frame construction method that best fits their budget and research needs.

## 2. Background and problem

This research was undertaken as a methodological supplement to the National Social Life and Health in Aging Project (NSHAP) using field listing and screening for the Health and retirement Survey (HRS), both NIH/NIA projects. [1]

Traditional listing is a method of address frame generation created by field staff, known as "listers", who record addresses present in defined geographies in a systematic manner (Kish 1965). Traditional listing has been considered the optimal or "gold standard" for frame construction; it is however relatively time consuming and extremely costly.

The present research was motivated by the desire to explore alternatives to time-consuming and expensive in-field listing of HUs for area-probability samples. It is part of a larger industry-wide research into list-based sampling frames. RTI began the process in 2001 with a direct (non-

[1] This research was supported by National Institute on Aging (NIA) grant R01-AG021487-02S1; L Waite, Principal Investigator; E Laumann, S Lindau, W Levinson, C O'Muircheartaigh, co-Investigators.

evaluative) use of USPS frames for an urban sample (reported in Iannacchione et al. 2003). NORC began its assessment of the USPS frame with a field evaluation of the frame in a subset of segments for the General Social Survey (GSS) in 2001-2, and continued with a series of field assessments in a set of inner-city surveys from 2002-2004.

NORC used the version of the USPS delivery sequence file (DSF) available from ADVO as the initial address frame for the entire country. All the addresses were geocoded and matched to Census geographies. There are a number of points at which error can be introduced, however. Each potential source of error raises fundamental questions about the utility of USPS-derived housing unit lists for area-probability sampling.

First, rural areas are at risk for undercoverage due to incomplete conversion to city-style addresses. Rural undercoverage is due to the prevalence of PO BOX and rural route addresses, which are valid for mailing but not useful for face-to-face interviewing. The presence of rural delivery types raise the risk of "multiplicity", by which a housing unit can be on a sampling frame more than once. For example, a housing unit could be represented by both a PO BOX as well as a physical dwelling.

Second, certain housing types are problematic. Temporary homes, such as trailer housing, are vulnerable to undercoverage. Similarly, newly constructed housing and recent conversions from non-residential to residential use take time to appear on the DSF. Informal or illegal apartments, such as attics, basements or space behind storefronts may also not appear on the DSF. Some types of apartment buildings also pose a risk of undercoverage, where mail is delivered without apartment numbers or where apartments do not have separate mailing addresses; these present a different type of challenge to the interviewer.

Third, there may be issues with geocoding the addresses and merging in the relevant geographic information. Specifically, TIGER base maps are known to be deficient in certain areas and subject to spatial error, making them off-set from streets. Because we use TIGER-derived files to geocode addresses and determine their census block, any spatial displacement will create error in address location. The analyst is also dependent on the synchronization of the address list, geocoding database, and TIGER area files. If any of these are out of synch, error will result.

Fourth, we are dependent on the timeliness with which the USPS lists are updated. If the update cycle is too coarse to consider new development or the razing of obsolete housing, we will have imperfect lists.

Despite these problems, many surveys have begun using list-based sampling frames. Several projects currently underway and completed by NORC, including the General Social Survey (2004, 2006, 2008) and the Survey of Consumer Finances (2004, 2007) have used USPS list-based frames for interviewing. The primary motivation in this switch to list based sampling has been the cost-savings. These surveys are willing to accept the risk of error from the sources described above in return for a substantial reduction in the cost of frame construction.

Of course, traditional listing is also subject to error. Ironically, research into the accuracy of the address lists has revealed shortcomings in traditional listing that had not been fully recognized. For example, traditional listing is known to have undercoverage in environments with multi unit buildings, coach houses, or poorly-labeled or informal housing units. In addition, the quality of traditional lists are largely influenced by the training and experience of available field staff.

No listing method is error-free. What survey designers need is guidance in how to select the best listing method for their project, given the sample design and budget. For this it is necessary to judge the potential for using USPS-derived lists a-priori. With advance knowledge of the quality of segments, we could rationally trade off accuracy against cost.

## 3. Methodology

Our research project began with the Health and Retirement Study's area-probability frame. In 2004 ISR listed and screened 549 segments. We consider all listings by HRS listers within the selected to comprise the T frame.

We had previously compiled a USPS address lists for all delivery points in the USA. We geocoded each address using the MapMarker Plus geocoding software package, to determine the longitude and latitude. Following geocoding we identified the geographic boundaries of all the segments in the sample, and used GDT Census blocks to represent them in MapInfo Professional GIS software.

We defined (i) the T frame as the listings for these segments provided by the field listers; (ii) the U frame as addresses that the geocoding program placed inside the segment boundaries; (iii) the U+ frame as all delivery points that geocoded inside the segments or within 300' of the segment boundaries; (iv) the USPS frame as all delivery points, no matter where they were located.

In 2005 NORC compared the addresses from the traditional ISR listing (T) to those on the USPS list (U) for the 549 segments in the HRS/NSHAP sample. The results were somewhat disappointing in that the overall match rates were considerably lower that expected (O'Muircheartaigh et. al 2005). There were strong indications that each of the approaches was unsuccessful in particular situations; the failure of traditional listing was particularly significant, given its dominance in national studies.

Consequently, we decided to undertake validation on a subsample of the 549 segments. We selected 100 segments for field verification/testing of the two lists; we conducted the field validation during the summer of 2005 as a supplement to the NSHAP field work.

The motivation behind this effort was to identify the level of agreement, to explore factors related to quality of each frame, and to estimate national coverage of the two frames. We checked (in the field) every address that was on either the traditional list (T) and/or the geocoded USPS (U). The result of the field verification is a new superior best or "B" frame.

The B frame constitutes the new gold standard because it contains all HUs identified as being in the segments. Thus, we have the following representations of the world: T, the traditionally-listed frame; U, the USPS addresses that geocode in selected segments; U+, the USPS addresses that geocode within a buffer of selected segments; USPS, the entire USPS list within the United States; B, the members of the U and T lists that were determined to actually be located with the segment by a secondary field effort. Figure 1 is a Venn diagram of the possible relationships among the frames.

Figure 1. Conceptual Diagram of Frames

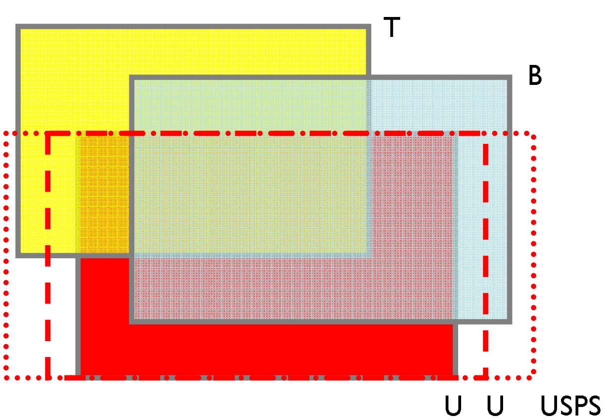


We then developed a predictive model using the SEARCH binary segmentation algorithm (Morgan and Sonquist 1963; Sonquist, Baker, and Morgan, 1974). SEARCH is an interactive detection program that divides a data set though a series of binary splits into a mutually exclusive series of subgroups (see: SEARCH manual). The splits are selected at each step to maximize the variance explained by splitting the candidate set into two groups. Such programs are used in marketing to identify parts of the market that behave differently (market segments or niches); here we are interested in identifying a model that will permit a priori identification of types of areas that will have different frame quality.

Even though we now consider it unlikely that field listing would be undertaken without the use of an address frame as a first stage, there are two alternatives that we should always consider. The first is to accept the list as it is without enhancement; the second is to build on the existing list in the field (list enhancement). Field enhancement of lists is quite expensive, and is impractical to implement quickly on a national scale (or even on a local scale for many surveys). Other alternatives, such as missed housing units procedures are ineffective and do not compensate for coverage issues satisfactorily.

For this project we have two aims in mind: the long term goal is to provide a predictive model that will enable us to tell in advance whether field enhancement will be necessary for a particular segment or class of segments; the more immediate goal is the develop an understanding of the factors that are most important in determining quality.

There is a wide variety of possibly relevant information available about a segment prior to fieldwork. For example, there are many Census products that can help describe the population in small area geographies, including race/ethnicity, income, education, etc. In addition, one can use Geographic Information Systems (GIS) to derive a wide range of physical measures, including population density, segment area, boundary types, the presence of water features, percent city-style addresses, etc.

Our key measure of quality is the proportion of the 'true' addresses contained in a candidate frame; this is the criterion variable we used in the SEARCH algorithm. In the figures and tables below, we define this as the intersection of the candidate frame with B: specifically, the percentage of the addresses in B correctly identified in the candidate frame. We examine in turn the intersections of the B frame with USPS, U, and T.

As predictor or explanatory variables, we used the following:

1. PSU selection stratum, as determined by city-style address concentration at the county or MSA level
2. Dummy variables classifying the holistic appearance of a segment as being urban, suburban, or rural
3. The presence of water features
4. The concentration of city-style addresses at the Census block level, as derived from TEA code
5. Percent housing units in urban areas from Census 2000 at the tract level
6. Percent housing units occupied at the tract level from Census 2000
7. The ratio of the U frame count to the Census 2000 housing unit count
8. A geospatial measure of the number of housing units on external streets that would be considered to be more subject to geocoding error than internal streets
9. Segment percent city-style by the USPS list at the ZIP code level
10. Percent White non-Latino from Census 2000 for the segment
11. Segment area, in $mi^2$
12. Segment block count
13. Segment population density, in $pop/mi^2$
14. Tract median income
15. Category of the Census Core Based Statistical Area (CBSA) e.g., metropolitan, micropolitan, or rural

## 4. Results and Discussion

These results are weighted to represent the whole of the US, based on the selection probabilities of the segments; our previous results were based on unweighted segment totals. National estimates are shown in table 1. We can see that the weighted estimates are more favorable to traditional listing than were the unweighted estimates; this is because rural segments (in which the T list performs better) had on average lower probabilities of selection.

If we exclude segments in which there were no street-style USPS addresses (and in which we could not use the address frame) the USPS-based frames perform significantly better. This exclusion essentially re-defines the target population. Thus, the estimates shown in table 1 should be considered to be perhaps unduly harsh as a reflection of the performance of USPS-based frames, while those in table 2 (that exclude five particularly problematic segments) are unduly generous.

*Table 1. Overall Intersections of B with U and T*

*% of B*

| *Frame* | *All Segments Unweighted* | *All Segments Weighted* |
|---|---|---|
| T | 73% | 80% |
| U | 77% | 74% |
| U+ | 83% | 79% |
| USPS | 84% | 81% |

*Table 2. Intersections of B with U and T (excluding 5 segments)*

*% of B*

| *Frame* | *National, Excluding Problem Segments, Weighted* | *Urban, Excluding Problem Segments, Weighted* |
|---|---|---|
| T | 79% | 75% |
| U | 79% | 86% |
| U+ | 84% | 89% |
| USPS | 86% | 90% |

The most informative presentation of the results of the SEARCH segmentation program is to show "dendrograms", or tree diagrams. For each node in the diagram two pieces of information are given: (i) the number of segments in that particular class; and (ii) the weighted estimate of the success rate (intersection with B) for that class. Thus, the top level in each dendrogram shows all 100 segments. When the algorithm determines that no further subdivisions can usefully be generated, this is called a "terminal" group. Terminal groups are shaded, while "pending" groups that can be further split are un-shaded. At each split the factor that was used by the model to make the partition is labeled. It is important to keep in mind that that there may be substantial instability in the sequence of splits, and alternative sequences may have been almost equally effective. Had other predictor variables been used, the splits might also have been quite different.

The dendrogram for the intersection of B in USPS is displayed in figure 2, showing at the top level the national estimate of B ∩ USPS of 81%, including those missing USPS lists entirely and subject to other structural deficiciencies. B ∩ USPS is first intuitively split on population density, with the more successful segments being of high to moderate population density (B ∩ USPS = 94%) against a group of lower population density (B ∩ USPS = 66%).

The high density segments are then split immediately into two terminal groups based on tract median income. There are 32 segments in the very effective 98% B ∩ USPS group with high to moderate median income. The other terminal group in the high-density stratum consists of 28 segments with a B ∩ USPS of 90% that is differentiated by lower median income. So, we can say that in general high population density, high-income segments have effective coverage of B with USPS.

The right side of the dendrogram begins with 40 segments of low to moderate population density. These 40 segments are first split on the ratio of the U frame count to the Census housing unit estimate. Such a measure is informative because if the U frame is considerably smaller than the Census, we can either suspect an area in transition (such as through demolition and rebuilding) or having PO BOXes that do not appear on the U.

The 28 segments of high-moderate U/Census are split on the percent of HUs that are city style, as derived from the Census TEA code. A terminal group is formed by 5 segments that have a low percentage of housing units in TEA 1 blocks, with an average B ∩ USPS of 62%. Such segments would probably require traditional listing, as they have a low ratio of U/Census and relatively few city-style addresses.

The 23 segments of high-moderate percent TEA1 are split into two terminal groups on population density, with high density segments being grouped in the very favorable 96% B ∩ USPS and lower density segments in the moderately successful 82% B ∩ USPS. Both terminal groups could be considered generally effective, however. While the dendrogram in figure 2 illustrating B ∩ USPS is generally intuitive with respect to the influence of physical factors, such as urbanicity or population density, we can observe some issues with the USPS list. The USPS list can be deficient in areas that are not obvious, as would be rural areas or those missing city-style addresses. Emerging or transitional parts of the USA, such as new suburbs on formerly agricultural land, represent a problem for the USPS list. In addition, urban areas with small multi-unit structures have been shown to be both subject to undercoverage as well as sensitive to geocoding and GIS database insufficiencies. We also see idiosyncrasies in postal delivery that complicate geocoding and prediction. Nonetheless, it is possible to generalize as to the categories of segments that can be expected to be effective with respect to the USPS list*:* dense areas, especially of moderate to high income, determined by the Census to have city-style addresses.

Because segments are often selected based on Census rather than postal geographies, it is necessary to consider the U in addition to the USPS frame. The difference between the U and USPS frames is that the former represents those addresses believed to be located within selected segments, and is so dependent on the quality of geocoding addresses to a given small area. We can therefore expect particular issues to influence the accuracy of the U list, including geocoding accuracy; postal irregularities; TIGER grid locations; and invisible block boundaries.

Figure 3 illustrates a predictive model for B ∩ U, which again uses only information available prior to fieldwork. We can see the national estimate of B ∩ U for all 100 segments was 74%, as shown in the top-level section.

As with the B ∩ ADVO model, the first split for B ∩ U is population density, with the left side being high density and high match and the right side moderate to low density and lower match. The left high-match side is split into two terminal groups by PSU selection stratum, which defines a county or metro-area in terms of the prevalence of city-style addresses. One of the two is a large group of 53 segments at a favorable 91% B ∩ U in strata 1,2 (large and medium urban areas). Not all segments in this group have high intersections with B ∩ U, however, due to localized geocoding irregularities; geocoding is an issue even in segments with favorable conditions.

The other terminal group on the left side of figure 3 contains seven segments in the rural parts of rural areas (Stratum 3) or the rural parts of large urban areas (stratum 4), with B ∩ U of 75%. These segments are often large and complex physically, as they can contain numerous Census blocks.

If we consider the right side of the top level of figure 3, these 40 segments are of moderate to low population density and are therefore generally larger and more rural than their counterparts on the left. We split the 40 lower density segments on a GIS-derived metric that quantifies the percentage of housing units that are on exposed (outer) streets, and therefore should be more sensitive to geocoding error than addresses on internal streets. In practice, this metric became a surrogate for urbanicity as smaller, denser urban segments have a larger share of housing units on outer streets. So, those with a low rating on this metric, and fewer housing units on exposed streets, in fact had a less favorable intersection between B and U.

We can then split the 14 segments with few households on exposed streets into two terminal groups by the ratio of U to the Census HU count. One terminal group contains six segments and has a moderate ratio of U to the Census, meaning U is comparable to the Census HU count. Segments in the group of six tend to have physical factors that negatively affect geocoding, such as urban blight or irregular street numbering, and an overall B ∩ U of 79%.

In addition to being of low population density, the other terminal group containing eight segments has a very low ratio of U to the Census. One would certainly choose to list segments in this group as the B ∩ U is only 10%, due to the prevalence of PO BOXes and the absence of city-style addresses.

The other group split from the 40 segments on the right side of figure 3 contained 26 segments having a large share of addresses on exposed streets. These 26 segments are then split on the percentage of city-style addresses in the segment's ZIP Code. Intuitively, the percent city-style positively influences the intersection between B and U.

Those segments that contain mostly city-style addresses are split on the left side by the ratio of U to the Census. Perhaps counter-intuitively, those where U/Census is very high perform considerably worse those where U/Census is moderate.

On the right side, we split the 15 segments with a moderate amount of city style addresses into two terminal groups based again on U/Census. One group contains 8 segments with a high ratio of U/Census and B ∩ U = 76%. The other has 7 segments with a low ratio of U/Census and B∩ U = 30%. This last group contains new suburbs that are rapidly growing as well as some geocoding problems that negatively affect the U frame.

The model in figure 3, while complex, does provide insight into the categories of segments that tend to be effectively represented by the U frame. We would expect dense segments with city-style addresses and a moderate ratio of U to the Census HU count to be successful. Segments with a very high or very low ratio of U to the Census HU count imply geocoding problems or the prevalence of PO BOXes, both of which lead to poor representation of B by U.

Our last analysis is shown in figure 4, a dendrogram of how B intersects the traditional or T list. On the surface, B ∩ T appears more convoluted than B ∩ U or B ∩ ADVO. We modeled B ∩ T using the same variables in the predictive model as the others, and saw a high degree of instability in the sequence of splits. Variation in the quality of T listing was therefore less systematic than B or USPS listing, probably due to the larger human component.

We can see that overall coverage with B and T is 80%, which is reasonable but lower than what might have been expected. The first split for B ∩ T is by income, with the lowest income quintile being separated from the upper four quartiles. Median income can be expected to influence factors that would challenge listers, and therefore affect B ∩ T, e.g., low income areas can be subject to trailer parks, multi-unit buildings, illegal apartments, etc.

The low-income group is then split by percent white non-Latino into high B ∩ T and low B ∩ T terminal groups. The more white non-Latino group has an overall worse B ∩ T than the less white non-Latino group, but that result is heavily affected by one segment with a very low B ∩ T. The more white non- Latino terminal group of six segments is heavily urban, and contains both an urban trailer park and new urban development. Conversely, the less white non-Latino terminal group is a mix of urban, suburban, and rural segments and generally has a moderate to high B ∩ T.

On the left side of figure 4, having higher income, we first split on the count of census blocks in a segment into a set of 56 with many blocks and a set of 24 with fewer. The set of 56 with more blocks is then split by physical area into a terminal group of 28 smaller segments (B ∩ T = 93%). This group of 28 small urban and suburban segments is the most successful overall with respect to T.

The non-terminal group with 24 segments, defined as having moderate to high income and a moderate number of blocks, is then split by urbanicity. The terminal urban group did well on all intersections, and had B ∩ T = 90%. The other half of the high-income set, also containing 28 segments, is split by the ratio of U to the Census into two intermediate groups. Ratio of U to the Census was demonstrated to influence how reality (B) is represented by U and USPS, but is also important to T as it indicates areas of rapid growth that would be difficult to list (and thus have a high ratio of U to the Census).

The group of 15 segments having a high to moderate ratio of U to the Census is split into two terminal groups by percent white non-Latino. One terminal group contains 7 segments having a high percent white non-Latino (and a high ratio of U /Census) with a mediocre B ∩ T of 63%. Challenging segments in this group tended to be suburban environments that would hinder listing, including cul-de-sacs and discontinuous blocks.

The other terminal group broken by white non-Latino contains eight segments that are less white non- Latino than the group with seven segments. These segments are often older 'inner ring' suburbs characterized by blight, irregular blocks/cul-de-sacs, or trailer housing.

Two terminal groups are split from the 13 segments that had a low ratio of U/Census, large areas, high block counts, and high incomes. The first consists of seven segments with relatively high incomes, and an overall B ∩ T of 92%. The other, containing six, is typified by irregular segment in urban or suburban environments with new construction. So, both terminal groups have problematic segments (large, irregular), and are typically better for T than they would be for U or ADVO.

Let us finally consider the group containing 24 segments split from the 80 high-income segments on the left side of figure 4. The group of 24 having lower block counts is first split on the left side into a non-terminal group of 12 non-urban segments. These 12 segments are then split into two terminal groups by PSU selection stratum. Those six segments in the urban stratum 1 were classified into in a terminal group of six segments in suburban environments. Lastly, there is a second terminal group of six segments in the less urban strata 2 and 3 that counter intuitively out-performed the group in stratum 1.

It is clear that there are a number of issues that can affect the effectiveness of traditional listing. Areas undergoing redevelopment or change appear to be difficult to list and therefore subject to undercoverage. Imperfections in the traditional listing are also widespread, however, and not just confined to particular environments or categories of segments. Some variation also appears to be idiosyncratic, with listers missing parts of segments for no systematic reason. Also, it should be emphasized that errors and undercoverage are not confined to one listing organization.

## 5. Conclusions

Our first conclusion is that while the USPS frame is not a panacea, it constitutes a breakthrough for small-scale low-cost surveys. For high quality national samples USPS samples alone are not adequate; geographic correspondence is a problem, as embodied by geocoding error. There are also implications for large-scale high quality surveys, such as the National Children's Study. Field listing has been shown to be imperfect, and so one should take an integrated approach to be successful.

Secondly, we contend that USPS-assisted samples are the future. One could start a project with USPS-derived lists. Then, it would be possible to upgrade the lists where possible from additional sources. Once would then do field enhancement, either complete, targeted, or via missed housing unit (MHU) procedures as determined by budget. Going forward, we will be doing more work on the *a priori* identification of problem areas.

## Acknowledgements

We appreciate help of the NSHAP NORC project team, especially Stephen Smith, Project Director, and Sara Zuckerbraun, Associate Project Director. In addition, we recognize the contributions of Heidi Upchurch and Katie Dekker.

## References

Iannacchione, V.G., Staab, J.M., & Redden, D.T. (2003). Evaluating the use of residential mailing lists in a metropolitan household survey. Public Opinion Quarterly, 67 (2):202-210.

Kish, Leslie. 1965. Survey Sampling. New York: John Wiley and Sons, Inc.

Morgan, James N. and John A. Sonquist (1963), "Problems in the Analysis of Survey Data, and a Proposal", Journal of the American Statistical Association, 58:415-435.

O'Muircheartaigh, Colm, Stephanie Eckman and Charlene Weiss. 2003. "Traditional and Enhanced Field Listing for Probability Sampling." 2002 Proceedings of the Section on Survey Research Methods, American Statistical Association.

O'Muircheartaigh, Colm, Stephanie Eckman, Ned English, James Lepkowski, and Steven Heeringa. *Comparison of Traditional Listings and USPS Address Database as a Frame for National Area Probability Samples.* Presented at American Association for Public Opinion Research Conference, May 2005, Miami Beach, Fl.

O'Muircheartaigh, Colm, Ned English, Stephanie Eckman, Heidi Upchurch, Erika Garcia, and James Lepkowski. 2006. "Validating a Sampling Revolution: Benchmarking Address Lists against Traditional Listing". 2006 Proceedings of the Section on Survey Research Methods, American Statistical Association.

Sonquist, J. A., E. L. Baker, and J. N. Morgan. *Searching for Structure*. Revised edition. Ann Arbor: Institute for Social Research, The University of Michigan, 1974.

Figure 2- Dendrogram of B ∩ USPS

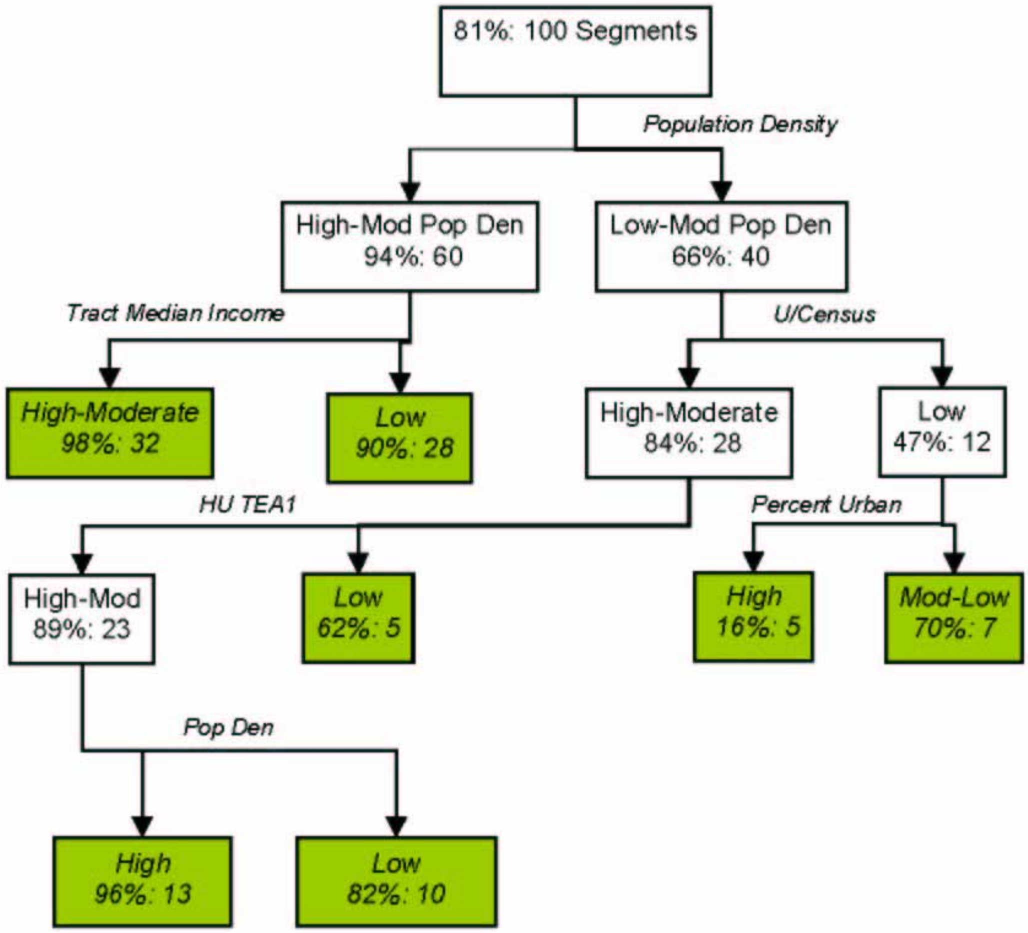

Figure 3- Dendrogram of B IN U

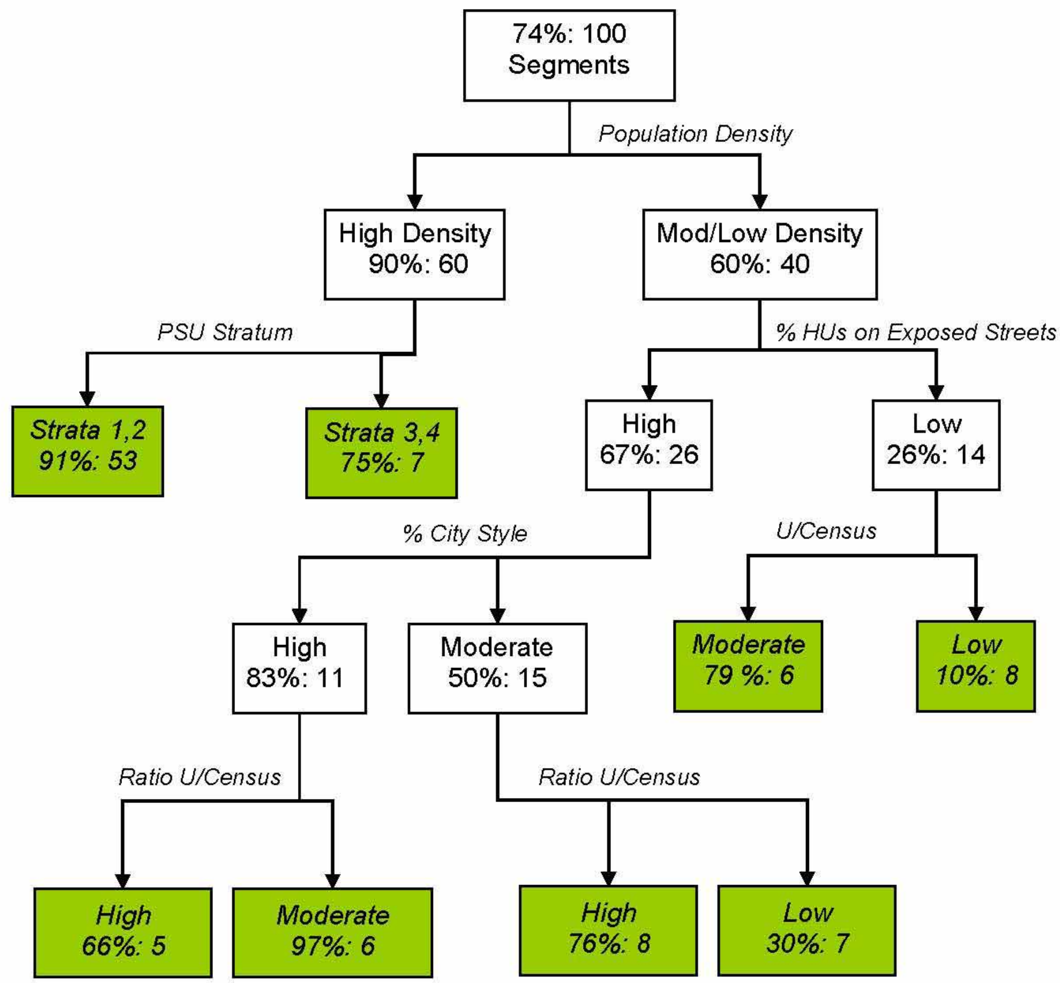

Figure 4- Dendrogram of B in T