

Classification of Address Register Coverage Rates - A *Field Study*

Gavin Thompson and Claude Turmelle, Statistics Canada
R.H. Coats Bldg., 16th floor, Tunney's Pasture, 120 Parkdale, Ottawa,
Ontario, Canada, K1A 0T6
Gavin.thompson@statcan.ca

Keywords: Labour Force Survey, Address Register, Coverage Rates.

Introduction

Every 10 years the Canadian Labour Force Survey (LFS) reviews its methodology to seek improvements in the manner with which the survey is conducted.

The LFS has a two staged design; first geographic clusters are selected, and within each selected cluster a sample of dwellings is taken. In order to sample in the cluster, the complete list of dwellings must be known. Traditionally, the list was created by interviewers touring the cluster and noting each dwelling they saw.

One potential innovation brought forward for review was the use of the Address Register (AR). Being a list of nearly 90% of all dwellings in Canada, the AR has the potential to both help reduce the amount of time spent locating dwellings and increase the accuracy with which dwellings are found.

In some areas the AR is accurate enough that no field work by the interviewer is needed. In other cases, however, an interviewer is still needed to validate the list from the AR.

In order to determine the best strategy to incorporate the use of the AR in the LFS, a field test was performed. There were two main objectives for this field test. The first objective was to validate earlier results showing that the AR had the overall quality necessary for the LFS. The second objective was to try and determine if a set of rules could be constructed which could separate clusters where the AR is accurate enough not to need interviewer intervention, from those clusters that do need interviewer intervention.

This paper will give a brief review of the Canadian LFS and the AR, report on the methodology and results of the field test, and discuss the applications of the test results to the LFS.

The Canadian Labour Force Survey

The Canadian Labour Force Survey is used to measure employment rates in various labour sectors of the Canadian economy. It produces a general 'unemployment rate' and also produces many other labour indicators including wage and union participation.

The survey has a complex two-staged design. The primary sampling units are geographic clusters; the secondary sampling units are dwellings. The clusters are regrouped into 6 rotation groups. Each month, all the selected dwellings of one rotation group rotate out and are replaced by another set of dwellings. This results in each cluster remaining in the survey for several years and dwellings remaining in the survey for 6 months. Interviews are conducted in approximately 53,000 dwellings each month.

The country is split into 58,266 clusters, with 6575 clusters selected for the survey. These clusters are designed to include roughly 150-200 dwellings. Clusters are also restricted in geographical area. Clusters may contain less than the 150 dwellings if it is required to keep the geographical size of the cluster manageable.

Because a sample of dwellings is being selected in each of the 6575 clusters, the complete list of dwellings must be known in each cluster. Currently, this requires each cluster to be 'listed' by an interviewer. The interviewer visits the cluster to construct an up to date list of all dwellings.

Listing Methods

The current method for listing all the dwellings in a selected cluster is to send an interviewer out 6 months before the inclusion of a cluster into the survey. The interviewer then tours the entire cluster and takes note of all the dwellings. In this paper, this will be referred to as the 'Current method' of listing a cluster.

The test introduced two new forms of listing. One is a method which exclusively used the AR, and was not verified by an interviewer, this will be called the 'AR

method'. This method did not require any work in the field and the test analysis only required that the address register list be compared to the final list of addresses compiled by the interviewers.

The second method examined was the 'Pre-list method'; in this method the interviewer was supplied with the list of dwellings in the cluster from the address register. The interviewer then validated this list by touring the entire cluster. This method helped the interviewer locate hidden dwellings, dwellings that may be inside or behind other dwellings for example. This method also cut down on the amount of data entry done by the interviewer.

As stated earlier, one of the goals of this test was to determine if a set of rules could be created so that when a cluster is added to the survey, it will be assigned to one of three groups:

Group 1: Clusters which have a good enough AR coverage that the AR method will provide a good frame from which to sample dwellings.

Group 2: Clusters which contain a good coverage rate, but will require an interviewer to use the pre-list method.

Group 3: Clusters where the AR offers little or no help to the interviewer.

The Address Register

The AR is a list of over 12.5 million Canadian dwellings. The majority of these dwellings are collected from census information, while the rest come from updates from administrative files.

The AR was originally designed to be maintained for communities with a population over 50,000. It was recently expanded to include less populated regions. Because of the history of maintenance for large communities the AR is more accurate in these communities. With time it will also become more accurate in less populated regions as well.

Because the census is only taken every 5 years (years ending in '1' or '6', most recently 2001) the quality of the AR can deteriorate over time. While most of the change will be from new construction, there are also other possibilities including businesses being converted to private dwellings (and vice versa) and single family dwellings being converted into apartments.

In spite of these two issues, the overall coverage rate of the AR is approximately 90%. With the knowledge that the coverage rates are fairly good, the field test is more concerned with the ability to predetermine which

clusters will have good coverage rates, and which will have coverage that will have to be verified.

In order to be a useful listing tool, the addresses from the AR must be linked to their exact geographic location. This linkage is done using the National Geographic Database. For an address to be properly located, the database must be able to determine the cross-streets that bound the address. This will determine the exact block-face that the address can be coded to. If an address can be coded to a block-face it is defined as 'structured', if an address can be coded to the cluster level, but not to the block-face level it is defined as 'unstructured'.

The Field Test

The field test was conducted in November 2003 in five cities. Ottawa, Toronto, Calgary and Vancouver were selected because they are considered large and significantly growing cities, and Trois-Rivières was selected as a medium sized city.

In total 180 clusters were used. In each of the major cities 40 clusters were used, and in Trois-Rivières 20 clusters were studied.

The field test took approximately 3 weeks to complete. In the first week half of the clusters in each city were randomly assigned to be listed with the current method of listing and half were assigned to the pre-list method. In the second week the methods are switched for each cluster. In the third week, the two methods were compared for accuracy.

When the test was designed, each interviewer was supposed to do five listings with each method and do some of each method in each of the two weeks. However, due to operational constraints, this was not always possible, although each interviewer did use both methods.

Once a cluster had been listed both ways, the two lists were then sent to a third interviewer. This interviewer examined the two lists and noted any discrepancies between the two methods. If no discrepancies were found, then each list was considered 'final', they both had the true list of addresses for the cluster. If the two lists did not agree, then the interviewer would go out to the field and verify each dwelling that was identified by one method and not the other. This third interviewer had the final say as to whether a dwelling existed or not. The final list constructed by the third interviewer would then be considered the 'true' list of addresses for the cluster.

Statistics Measured

Six different statistics were measured for each cluster. For each of the three listing methods the following statistics were calculated:

Overcoverage: This is the number of dwellings that appear on a list that did not appear on the true list of addresses. These addresses were erroneously attributed to the cluster.

Undercoverage: This is the number of dwellings that fail to appear on a list that were on the true list of addresses. These are addresses that were erroneously excluded from the cluster.

General Results

The first part of the test was to determine the overall quality of the AR and Pre-List methods. If either of these methods returned with poor results, it would have put into question our ability to use them. It would have required that the set of rules be very precise about which group clusters should be split into, if the two methods were to be used at all.

As expected, the results came back very favourably. The AR has been tested for other sections of Statistics Canada, and we expected from their results that these results should be comparable. The primary focus of this test was to determine if the split of the clusters could be pre-determined. The design of the test produced the validation of the AR and Pre-List methods at no additional cost.

Table 1 shows the rates of undercoverage for the AR and Pre-List methods, as well as for the current listing method, for the five selected cities.

City	AR Method	Pre-List Method	Current Method
Trois-Rivières	4.1	2.5	7.6
Ottawa	4.6	0.7	3.9
Toronto	1.3	0.5	2.2
Calgary	6.0	0.3	3.2
Vancouver	5.4	3.1	6.0
Average	3.3	1.1	3.5

The undercoverage rates for both the AR method and the Pre-List method were better, on average, than the Current method. These results were expected, since the AR contains dwellings that may be difficult for the interviewer to find without the aid of the AR.

The AR method and Pre-List method both contain many more cases of dwellings which do not exist.

City	AR Method	Pre-List Method	Current Method
Trois-Rivières	9.7	6.5	2.5
Ottawa	4.4	0.6	0.8
Toronto	11.3	3.6	2.3
Calgary	8.1	0.01	0.9
Vancouver	5.6	2.6	1.8
Average	8.8	2.7	1.8

It is apparent from Table 2 that there is a high level of inter-city variability in the overcoverage rates. The same can be seen in Table 1, but to a lesser degree. We believe that this is primarily due to interviewer bias rather than a true difference in the amount of overcoverage and undercoverage in each city. It appears that the third interviewers in Calgary sided almost exclusively with the Pre-List method, whereas in some cities it seemed like there was some bias to side with the Current method. The interviewers were given a set of procedures which were slightly different from those they were used to; this resulted in the procedures being interpreted differently in different cities.

The Rules

In addition to concerns about frame quality, it was also essential that the new method be easy for the interviewers to use. If the new method is very difficult or confusing, any gains from the new system would potentially be lost through interviewer error. Unstructured addresses appear at the end of an interviewers list. They also appear alphabetically, rather than in the order they would appear on a street. Therefore, it became important to try and minimize the number of unstructured addresses in Groups 1 and 2. A high number of unstructured addresses will make Pre-List method listing much more difficult than the Current method.

There are two competing forces at work: the desire to maximize the number of clusters not requiring checking and the desire to minimize errors in the frame. Within the desire to minimize frame error there are two issues; the quality of the AR for the cluster and the ease with which the interviews can use the AR for the cluster. Both of these concerns were taken into account when constructing the rules.

Two sets of variables are used to construct the rules. The first set of variables relate to the nature of unstructured addresses in the cluster (unstructured road

count, unstructured address count and ratio unstructured). The second set of variables relates to the quality of the Address Register; they are coverage and growth variables (Address Register Coverage, Apartments size 2/3 ratio, Address Register Growth, Ratio of AR Growth). Below is the definition of each variable.

Unstructured Road Count – This is the number of roads containing only unstructured addresses. This is important for two reasons: first, a high number of unstructured streets may indicate a problem with the way the sequencing (order) of addresses for the cluster was done, and secondly, if a cluster contains many unstructured addresses the printed list is not in a very good order and is difficult for the interviewer to use.

Unstructured Address Count – This is the actual civic number count and not the dwelling count. We reasoned that having 40 single unit dwellings that are unstructured are more difficult for the interviewers to work with than having one unstructured building with 40 apartments.

Ratio Unstructured – This is the ratio of the number of unstructured dwellings in the cluster to the number of dwellings on the Address Register for that cluster.

Address Register Coverage – The theoretic coverage rate of the cluster can be calculated for each cluster. The coverage rate is the number of dwellings that appear on the AR and were considered valid in 2001, divided by the number of dwellings counted during the 2001 Census.

There are several reasons why coverage may not be 100%. It is possible that not all the dwellings in a cluster have been ‘captured’ on the AR. Some are incorrectly geocoded either into or out of the cluster (overcoverage and undercoverage respectively) and others may have been incorrectly added by an update

file; it may be that they are not a dwelling (for example a business) or in some cases do not exist at all.

This value is not the true undercoverage or overcoverage of a cluster. To calculate these values, it is necessary to know which values on the AR are false (overcoverage) and which dwellings in the cluster were missed (undercoverage). This value is an indication of how well the AR should perform in the cluster.

Apartments size 2/3 ratio – This variable is the ratio of the number of 2 or 3 unit addresses in the cluster. This was identified as a variable which may be important to include in the rules. Apartments are more likely to be erroneously put on the AR than a single dwelling unit. If there was a high number of units of size 2 or 3 verification becomes necessary to ensure the quality of the frame.

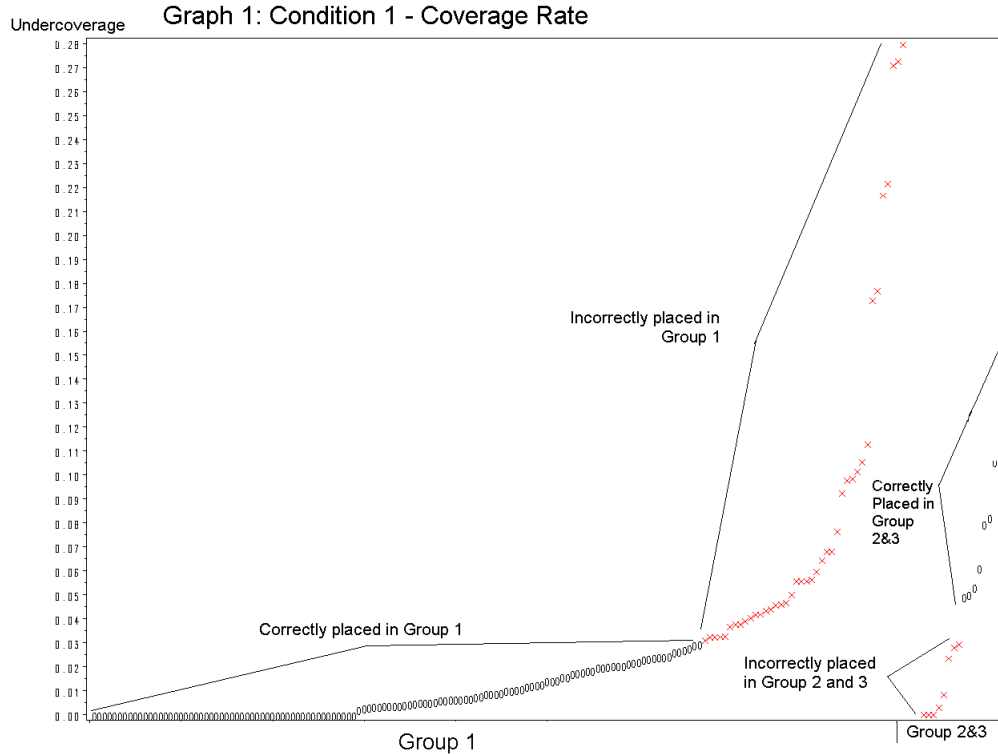
AR Growth – The growth of a cluster is measured by the number of dwellings on the AR which were added from administrative files since the last Census.

Ratio of AR Growth – This is the ratio of the growth in the cluster to the total number of dwellings in the cluster.

The way the rules separate clusters are as follows. If a cluster fails to pass a basic quality test, it is assigned to Group 3. If a cluster passes the basic quality test it is tested to see if it belongs in Group 1. If it does not belong in Group 1 it is then placed in Group 2.

A cluster is assigned to Group 3 if it fails any of the 3 following criteria:

- 1) If the **Unstructured Road Count** is greater than 4.
- 2) If the **Unstructured Road Count** is equal to 4 and the **Unstructured Address Count** is greater than 20.



- 3) If the **Address Register Coverage** is less than 90%.

In total, 29 clusters (or 16.1% of the test) failed this basic test and were assigned to Group 3.

There are 3 criteria for inclusion in Group 1:

- 1) The **Address Register Coverage** must be between 98% and 102%
- 2) The **Ratio Unstructured** has to be less than .1 or less than 0.2 if the **Unstructured Road Count** is only 1 and the **Unstructured Address Count** is less than 5.
- 3) The **Apartment Size 2/3 Ratio** is less than 0.4 and the **Address Register Growth** is less than 11 or the **Ratio of Address Register Growth** is less than .02; or the **Apartment Size 2/3 Ratio** is less than 0.2 and the **Ratio of Address Register Growth** is less than 0.035.

If a cluster fails any of the three conditions for Group 1, and did not end up in Group 3, then the cluster is assigned to Group 2. In total, 86 clusters (or 47.8% of the test) were assigned to Group 1 and 65 clusters (36.1%) were assigned to Group 2.

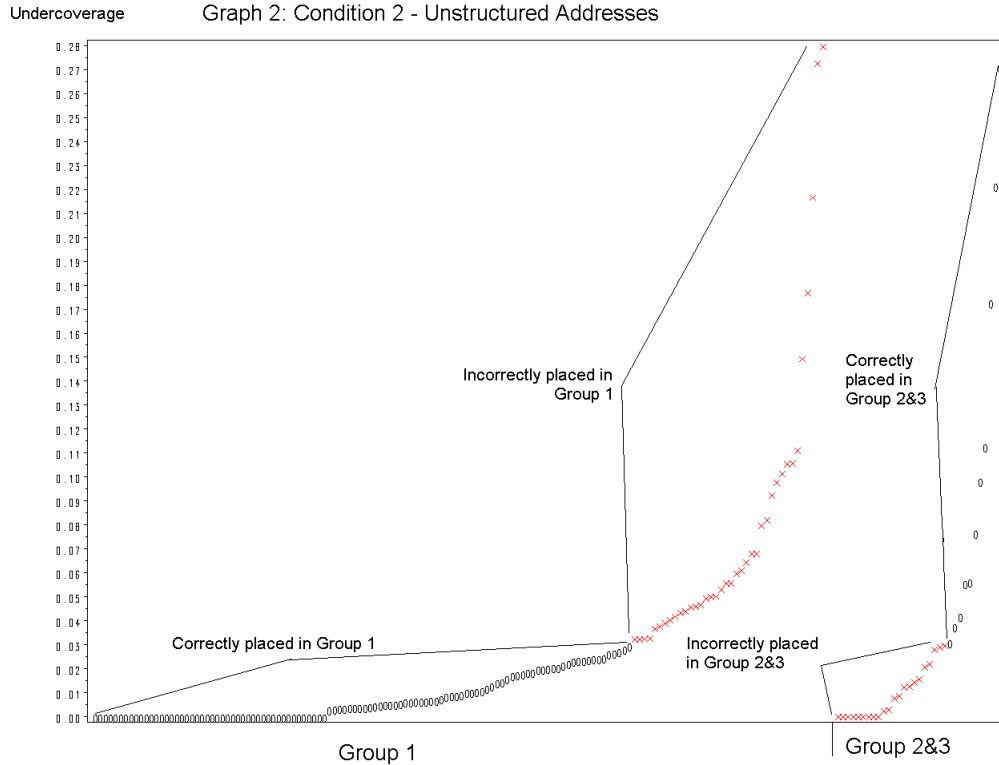
Results

The next section will review the results for the effect on undercoverage for the rules. The results for the overcoverage were very similar to the results for the undercoverage.

A level of 3% true undercoverage was used as the 'golden rule' for 'correctly' placing clusters in Group 1 or Groups 2 and 3. If a cluster had 3% or less AR method undercoverage, then we hypothesized it should be assigned to Group 1. If it had more than 3% undercoverage then it should be assigned to a group that would later be validated by an interviewer. This isn't the only criterion for a cluster being 'good', as an example, it is also important that the number of unstructured address be low.

Any cluster incorrectly assigned to Group 1 could result in a cluster with low AR quality receiving no validation from an interviewer. If a cluster is incorrectly assigned to Groups 2 or 3 it indicates that a good quality cluster will be verified by an interviewer. This results in extra work being done for relatively little (or no) gain.

It should be noted that each criterion, when taken separately, has a high number of clusters incorrectly assigned to Group 1, but when taken together there are much fewer clusters incorrectly assigned to Group 1.



Graph 1 shows the results for the first condition, that the **Address Register Coverage** must be between 98% and 102%. If this condition alone was used, 163 clusters would be assigned to Group 1 and 17 would be assigned to Groups 2 and 3. In total 120 would have been correctly assigned in Group 1 and 9 correctly assigned in Groups 2 and 3. 43 Clusters would have been incorrectly assigned to Group 1 and 8 would have been incorrectly assigned to Groups 2 and 3.

This criterion does not separate many clusters into Groups 2 and 3. This is true for the test, where the clusters only come from five cities.

The intention of this rule is to be able to apply them to the entire LFS, which contains many rural clusters with much lower coverage rates. The reason that the coverage rate is capped at 102% is that a level higher than that indicates a higher incidence of overcoverage.

Graph 2 shows that there are a very high number of clusters being incorrectly assigned to Group 1. This is not overly concerning, since all three criteria will be used at the same time.

Graph 3 shows that the criterion that includes the ratio of apartments of size 2 and 3 is better suited to separate clusters into Group 1 or Group 2 and 3. In total 125 clusters are assigned to Group 1, with 83.2% (104) being correctly assigned to Group 1. Not surprisingly, it

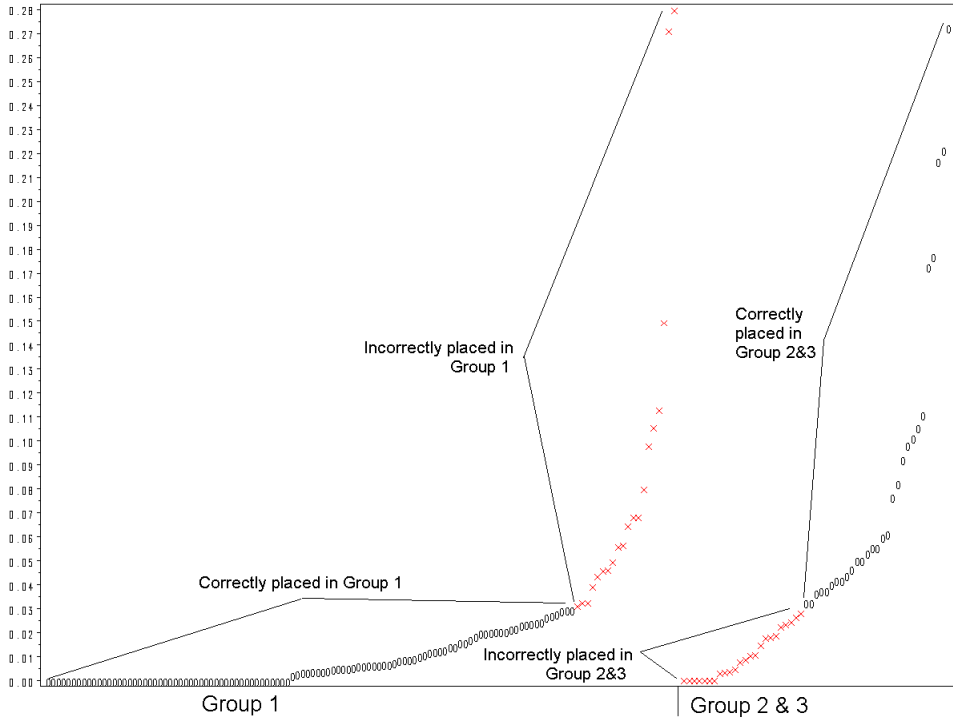
also minimizes the number of clusters incorrectly assigned to Group 2 and 3 (24).

The final group allocation can be seen Graph 4. The set of rules allocated 75 clusters correctly to Group 1, 29 were correctly allocated to Group 2 and 12 were correctly allocated to Group 3. In total 64 clusters were incorrectly allocated. 11 were incorrectly allocated to Group 1, 36 were incorrectly allocated to Group 2, and 17 were incorrectly allocated to Group 3.

Although the Group 3 clusters have been listed as 'correctly' or 'incorrectly' allocated as per the 3% Address Register undercoverage rule, the rule does not completely apply to Group 3 because these clusters can cause operational difficulty for the interviewer trying to use the Pre-List method. The Address Register does not work for the cluster very well and this will be reflected in the list that the interviewer would be working from.

Even with a few clusters that can be considered incorrectly allocated, the rules do very well at classifying the clusters into groups. Table 3 shows the undercoverage and Table 4 shows the overcoverage of the Address Register and the Method Specific coverage rates. 'Method Specific' is using the AR method for Group 1, Pre-List Method for Group 2 and the Current Method for Group 3.

Undercoverage Graph 3: Condition 3 - Apt 2&3 and AR Growth



Undercoverage Graph 4: Final Groupings

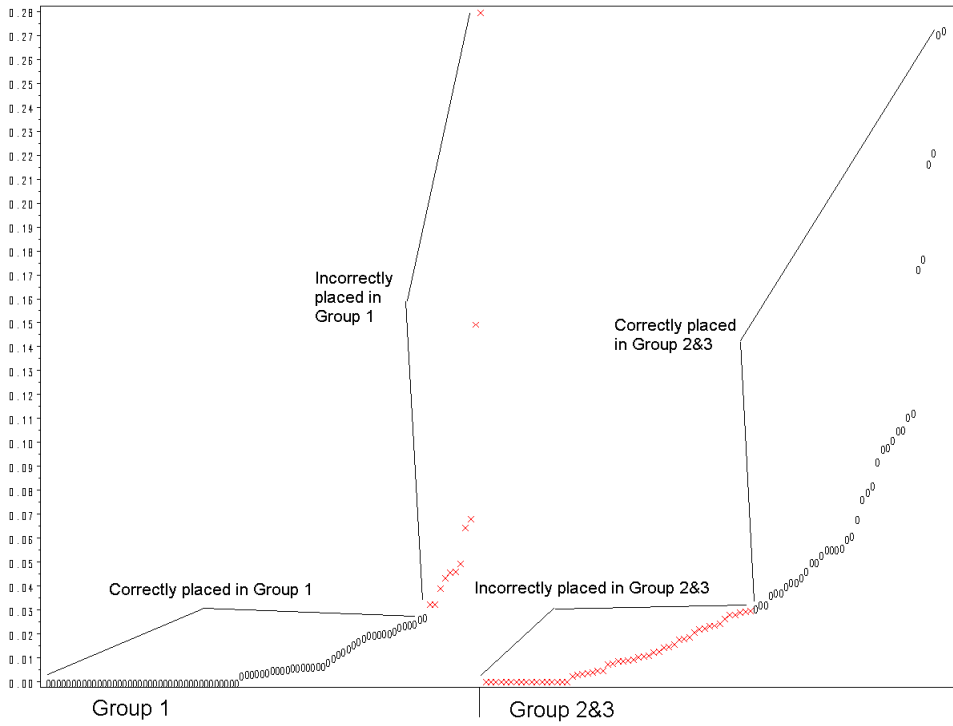


Table 3: Undercoverage rates (%) by Group

Group	AR Method	Pre-List Method	Current Method	Method Specific
1	1.2	1.2	2.0	1.2
2	6.4	2.8	6.2	2.8
3	5.3	0.6	4.3	4.3
Average	3.3	1.1	3.5	2.2

From the relatively low undercoverage rate from the AR in Group 1 compared to Group 2 or 3 we see that the rules split the groups effectively. It also shows that the AR method creates less undercoverage than the Current method does. A question about using the Pre-List method in Group 3 arises however, since the undercoverage is much lower. The Pre-List method is not used in Group 3 because of operational problems with the clusters. The Group 3 clusters in the rest of the country will not yield the same good results as Group 3 from the selected 5 cities.

Table 4: Overcoverage rates (%) by Group

Group	AR Method	Pre-List Method	Current Method	Method Specific
1	4.8	0.8	2.3	4.8
2	13.5	4.1	1.2	4.1
3	14.2	6.1	1.3	1.3
Average	8.8	2.7	1.8	4.0

Table 4 shows that the rules again split well between Group 1 and Groups 2 and 3. Although the new method has a higher rate of overcoverage, it is not considered too high.

Application to the Labour Force Survey

Although the test was conducted on a fairly homogeneous set of clusters (all coming from the urban parts of major cities), the results are promising enough to proceed with the application of this new method to the entire LFS.

Some of the rules may have to be changed to conform to the realities of rural clusters and there may be some new constraints that need to be added due to the potentially large size of some clusters.