SSC Annual Meeting, June 1996
Proceedings of the Survey Methods Section

MONTH-IN-SAMPLE EFFECTS FOR THE CANADIAN LABOUR FORCE SURVEY

F. Brisebois and H. Mantel[1]

ABSTRACT

The Canadian Labour Force Survey (LFS) uses a rotating panel design in which each sample dwelling remains in the sample for six consecutive months before being replaced by a new sample dwelling; thus one sixth of the sample is replaced each month. The question that we consider here is whether the LFS experiences month-in-sample bias, *i.e.*, whether the estimate of a characteristic based on the panel of dwellings which are in the sample for the *k*-th month differs significantly from the overall estimate. We also investigate whether the effect, if it exists, is consistent over regions of the country. Two approaches are considered: a modelling approach based on state space modelling of the time series of panel estimates, and the use of indices based on a comparison of panel estimates over the six month lifetime of a panel in the sample. We describe the methodology used and give a summary of the results for the two characteristics "unemployment rate" and "employment rate".

KEY WORDS: Non-response bias; Response bias; Rotating panel design.

RÉSUMÉ

L'enquête sur la population active du Canada (EPA) fait usage d'un plan d'échantillonnage par panel avec renouvellement de l'échantillon dans lequel chaque logement sélectionné reste dans l'échantillon pour six mois consécutifs avant d'être remplacé par un nouveau logement, ainsi un sixième de l'échantillon est remplacé chaque mois. La question que nous posons ici est de savoir si l'EPA souffre de biais dû au mois d'inclusion dans l'échantillon, c.-à-d., de savoir si l'estimation d'une caractéristique qui est basée sur le panel des logements qui sont dans l'échantillon pour le *k*-ième mois diffère de façon significative de l'estimation globale. De plus, nous vérifions si cet effet, s'il existe, est cohérent à travers les régions du pays. Deux approches sont considérées: une approche de modelage basée sur une modélisation d'espace d'états des séries chronologiques des estimations du panel, et l'utilisation d'indices basés sur une comparaison des estimations du panel sur la durée de vie de six mois d'un panel dans l'échantillon. Nous décrivons la méthodologie utilisée et donnons un sommaire des résultats pour les deux caractéristiques "taux de chômage" et "taux d'emploi".

MOTS CLÉS: Biais de non-réponse; biais de réponse; plan d'échantillonnage par panel avec renouvellement.

1. INTRODUCTION

The Canadian Labour Force Survey (LFS) is a monthly household survey which uses a rotating panel survey design in which each sample dwelling remains in the sample for six consecutive months before being replaced by a new sample dwelling; thus one sixth of the sample (one panel) is replaced each month.

A question of considerable interest is whether there are systematic differences among the different panel estimates which can be attributed to the number of months that the panel dwellings were in the sample. Such differences can arise because of systematic differences in non-response patterns for the different month-in-sample groups, or because of different response error propensities among the different groups.

For the LFS it is already known (Kennedy, Drew and Lorenz, 1994) that non-response rates vary among the month-in-sample groups, with the first month-in-sample group generally having the highest non-response rates, the second month-in-sample group having the lowest non-response, and non-response generally increasing from months in sample two through six. The high non-response in the first month is due to difficulties in contacting the household for the first time, particularly for households of size one or two. The increasing non-response rates for months two through six are due to increasing rates of refusal. Thus

[1] François Brisebois and Harold Mantel, Household Survey Methods Division, Statistics Canada, Tunney's Pasture, Ottawa, Ontario, Canada, K1A 0T6.

significant month-in-sample effects exist for LFS panel estimates of totals.

The question we investigate in this paper is whether or not month-in-sample effects exist also for LFS panel estimates of the employment and unemployment rates. If month-in-sample effects do exist for these variables, how large are they and is there a common pattern of month-in-sample effects for different regions of the country?

We analyse monthly LFS data for the 52 months from May 1990 to August 1994 and for five regions: Atlantic, Quebec, Ontario, Prairies, and British Columbia. Two approaches are considered: a modelling approach based on state space modelling of the time series of panel estimates, and the use of indices based on a comparison of panel estimates over the six month lifetime of a panel in the sample. These approaches are described in Sections 2 and 3 respectively. Results of the analyses are given in Section 4, and some discussion is presented in Section 5.

As mentioned above, we consider the variables employment rate and unemployment rate. Persons in the LFS population (*i.e.*, the civilian non-institutional population 15 years of age or older) are classified as being either employed, unemployed, or not in the labour force. Letting E denote the number of employed persons, U the number of unemployed, and NLF the number not in the labour force, the employment rate is defined as

$$ER = E/(E+U+NLF) \tag{1.1}$$

and the unemployment rate is defined as

$$UR = U/(E+U). \tag{1.2}$$

In what follows $ER_{k,t}$ denotes the subweighted panel estimate of ER and $UR_{k,t}$ denotes the subweighted panel estimate of UR for the panel that is in the sample for the k-th occasion at time t. These subweighted estimates are obtained from (1.1) and (1.2), respectively, by replacing E, U, and NLF by expansion estimates based on subweights. The subweights are derived from the raw survey weights by applying an adjustment for household non-response. We decided to use sub-weights instead of final weights, which are calibrated to various demographic characteristics such as age-sex group populations within provinces, since we thought that the calibration might obscure or change the differences among the month-in-sample groups making them harder to understand. The non-response adjustment was done within predefined geographic areas for all rotation groups simultaneously, *i.e.*, the non-response adjustment factors were the same for all six rotation groups. It would have been preferable to use subweights based on non-response adjustment within rotation groups; however, these were not available.

2. STATE SPACE MODEL

In order to study the month-in-sample effects for employment and unemployment rates we developed state space models for the time series of subweighted panel estimates $\{ER_{k,t}, k=1,...,6, t=1,...,52\}$ and $\{UR_{k,t}, k=1,...,6, t=1,...,52\}$. These models incorporate parameters which correspond to month-in-sample effects which allows for testing of interesting hypotheses and examination of patterns in the estimated effects. The approach here is modelled after that of Pfeffermann and Bleuer (1993).

Letting $\hat{r}_{k,t}$ denote either $ER_{k,t}$ or $UR_{k,t}$ our model has the following basic form:

$$\hat{r}_{k,t} = r_t + \beta_k + e_{k,t} \tag{2.1}$$

where $r_t + \beta_k$ is the expected value of $\hat{r}_{k,t}$, β_k is a month-in-sample effect (we arbitrarily set $\Sigma\beta_k = 0$), and $e_{k,t}$ is a sampling error term.

The rate r_t is assumed to follow a simple basic structural model over time:

$$\begin{aligned}
r_t &= L_t + S_t + I_t & & \\
L_t &= L_{t-1} + R_t + \eta_{L,t} & \eta_{L,t} &\sim (0,\sigma_L^2) \\
R_t &= R_{t-1} + \eta_{R,t} & \eta_{R,t} &\sim (0,\sigma_R^2) \\
S_t &= -\Sigma_{j=1}^{11} S_{t-j} + \eta_{S,t} & \eta_{S,t} &\sim (0,\sigma_S^2) \\
I_t &= \eta_{I,t} & \eta_{I,t} &\sim (0,\sigma_I^2)
\end{aligned} \tag{2.2}$$

where the innovations $\eta_{L,t}$, $\eta_{R,t}$, $\eta_{S,t}$, and $\eta_{I,t}$ are all mutually independent and also independent over time.

The month-in-sample effects β_k are assumed to be constant over time. This assumption was tested by comparing estimates of the β_ks based on the first few months of data to those based on the entire 52 months of data; no significant differences were found.

Finally, the sampling error terms are also correlated over time since 5/6 of the sample is retained from one month to the next and even for the panel that rotates there is some correlation since it is generally only the sample dwellings (the ultimate sampling units) which rotate while higher order sample units are retained. We first defined a new sampling error $\varepsilon_{k,t}$ as a multiple of $e_{k,t}$ to adjust for changes in the design effect over time due to changes in sample size or changes in the underlying rate; specifically:

$$\varepsilon_{k,t} = \left\{ n_{k,t} / \hat{r}_t (1 - \hat{r}_t) \right\}^{1/2} e_{k,t} \qquad (2.3)$$

where $n_{k,t}$ is the number of sample dwellings in the panel that is in the sample for the k-th month at time t, and $\hat{r}_t$ is the simple average of $\hat{r}_{k,t}$, $k = 1,\ldots,6$.

Following Pfeffermann and Bleuer (1993), we considered autoregressive models for $\varepsilon_{k,t}$. However, unlike them, we found that an AR(1) model was adequate; we tested this informally by looking at the empirical correlations of the innovations in the sampling error as estimated under the AR(1) model with the lag 2 estimated sampling errors. The model for $\varepsilon_{k,t}$ is then:

$$\varepsilon_{k,t} = \begin{cases} \rho_1 \varepsilon_{k-1,t-1} + \xi_{k,t} & \xi_{k,t} \sim (0, (1-\rho_1^2)\sigma_\varepsilon^2) \quad k = 2,\ldots,6 \\ \rho_2 \varepsilon_{6,t-1} + \xi_{1,t} & \xi_{1,t} \sim (0, (1-\rho_2^2)\sigma_\varepsilon^2) \quad k = 1 \end{cases} \qquad (2.4)$$

The overall model has seven hyper-parameters: σ_L^2, σ_R^2, σ_S^2, σ_I^2, σ_ε^2, ρ_1 and ρ_2. Given values for these hyper-parameters and a suitable initialization procedure the Kalman filter may be used to calculate best linear predictors of the quantities L_t, R_t, S_t, I_t, β_k and $\varepsilon_{k,t}$, and hence to find the probability density of the observed data given the proposed values of the hyper-parameters (assuming a normal distribution for the ηs in (2.2) and the ξs in (2.4)). To initialize the Kalman filter we set the estimate of the state vector for time t=0 to zeros and we assumed the mean squared error matrix of this estimate to be diagonal with σ_I^2 in the position corresponding to I_0, σ_ε^2 in the positions corresponding to $\varepsilon_{k,0}$, k=1,...,6, and 1 (*i.e.*, a large number) in the other positions. We then used a numerical procedure to find the maximum likelihood estimates of the hyper-parameters.

Likelihood ratio tests and a backwards elimination procedure were used to drop non-significant hyper-parameters from the model. Note that if the hyper-parameter σ_S^2 is found to be non-significant, this does not imply that the seasonal effect does not exist; it implies only that the seasonal effect is constant over time. Similarly, if σ_R^2 is found to be non-significant this implies only that the rate of change in the level L_t has a constant mean over time. The month-in-sample effects β_k were estimated in the final model and interesting hypotheses concerning these parameters were tested using model based procedures.

3. MONTH-IN-SAMPLE INDICES

As a verification of the results obtained from the model based approach, we also analysed the month-in-sample effects using month-in-sample indices. The basic idea of these indices is to compare estimates based on data from each month-in-sample group to the average over the six month-in-sample groups. As in Section 2 we let $\hat{r}_{k,t}$ denote either $ER_{k,t}$ or $UR_{k,t}$.

Bailar (1975) defined a month-in-sample k at time t index as

$$\tilde{I}_{k,t} = 6\hat{r}_{k,t} \Big/ \Sigma_{i=1}^{6} \hat{r}_{i,t}.$$

Since differences among the $\hat{r}_{k,t}$ for different k include sampling errors as well as month-in-sample effects, this index may be quite unstable.

Mohl (1991) defined a month-in-sample index as

$$\tilde{I}_{k,t} = 6\hat{r}_{k,t} \Big/ \Sigma_{i=1}^{6} \hat{r}_{i,t-k+i}.$$

Sampling errors are not of concern for this index since all of the estimated rates are based on the same panel and hence are based on the same sample; however, this index would be confounded with trends in the underlying true rate over time.

For our study we used a modified month-in-sample index that takes advantage of the strengths of both the Bailar and Mohl indices. As in the Mohl index, we base our index on six panel estimates from the same panel and thus we avoid confounding with sampling error; however, we first adjust these panel estimates to account for trends in the underlying rate. To be specific, we take

$$\begin{aligned} I_{k,t} &= 6 r^*_{k,t} \Big/ \Sigma_{i=1}^{6} r^*_{i,t-k+i}, \text{ where} \\ r^*_{i,t} &= \hat{r}_{i,t} - (\hat{r}_t - \Sigma_{j=1}^{6} \hat{r}_{t-i+j}/6) \end{aligned} \qquad (3.1)$$

and $\hat{r}_t$ is the simple average of $\hat{r}_{i,t}$, $i = 1,\ldots,6$. This index may be considered as a modification of the Mohl index with adjustments to account for trends in the underlying rate.

Another possibility, which we did not consider, is to adjust the Bailar index to account for differences among the sampling error effects. For example, we could have defined

$$\begin{aligned} \tilde{I}_{k,t} &= 6 r^\diamond_{k,t} \Big/ \Sigma_{i=1}^{6} r^\diamond_{i,t}, \text{ where} \\ r^\diamond_{i,t} &= \hat{r}_{i,t} - (\tilde{r}_{t-i} - \Sigma_{j=1}^{6} \hat{r}_{t-i+j}/6) \end{aligned}$$

and $\tilde{r}_t$ is the simple average of $\hat{r}_{i,t+i}$, $i = 1,...,6$. Here $\overset{\diamond}{r}_{i,t}$ is just $\hat{r}_{i,t}$ adjusted by a crude estimate of the sampling error for the panel which is in the sample for the i-th occasion at time t.

The 52 months of LFS data that we have contain 47 complete six month panel histories, so for each k the index in (3.1) can be calculated for 47 different values of t. Assuming that the expected value of $I_{k,t}$ is common for all t, call it I_k, we estimate I_k by the simple average of $I_{k,t}$ over t and use a t-test to test whether I_k is significantly different from 1.

We also calculated averages over time for the Bailar and Mohl indices. The results for the Bailar indices were very similar to those obtained for our modified index given in (3.1), which is not surprising since in averaging the Bailar indices over time the sampling errors would also be averaged. The results for the Mohl index were quite different, due to trends in the employment and unemployment rates.

4. RESULTS

4.1. State Space Model

Within each region likelihood ratio tests and a backwards elimination procedure were used to drop non-significant hyper-parameters from the model. Table 4.1 gives the final models obtained for the five regions for the variable "employment rate" and Table 4.2 gives the same for the variable "unemployment rate". Only two hyper-parameters, σ_ε^2 and ρ_1, were found to be significant for both variables and for all five regions. Note that the estimates of these hyper-parameters do not vary very much across regions. Note

Table 4.1 Final models and hyper-parameter estimates for the employment rate

	Atl.	Que.	Ont.	Pra.	B.C.
σ_ε^2	1.39	1.90	1.19	0.96	1.29
$\sigma_L^2 \times 10^8$	--	--	--	268	--
$\sigma_R^2 \times 10^8$	--	4.81	44.1	--	--
$\sigma_S^2 \times 10^8$	238	--	--	--	--
$\sigma_I^2 \times 10^8$	--	--	--	--	--
ρ_1	0.93	0.94	0.92	0.89	0.91
ρ_2	0.57	0.38	--	--	--

-- = not significant

Table 4.2 Final models and hyper-parameter estimates for the unemployment rate

	Atl.	Que.	Ont.	Pra.	B.C.
σ_ε^2	1.47	1.89	1.54	1.56	1.21
$\sigma_L^2 \times 10^8$	--	1130	--	341	544
$\sigma_R^2 \times 10^8$	2.02	--	101	--	--
$\sigma_S^2 \times 10^8$	--	--	--	--	--
$\sigma_I^2 \times 10^8$	--	--	--	--	--
ρ_1	0.67	0.71	0.63	0.73	0.53
ρ_2	--	--	--	0.33	--

-- = not significant

also that the ρ_1 parameters for the variable "employment rate" are much larger than those for "unemployment rate". This may be because the state of employment is generally more stable than the state of unemployment. The hyper-parameter σ_I^2 was never significant.

We also tried fitting the model to all five regions simultaneously using common values for the hyper-parameters; however, the likelihood ratio tests indicated that there was overwhelming evidence against such common hyper-parameters.

The employment rate series for the Atlantic region was the only series for which the parameter σ_S^2 was significant; nevertheless, there were significant, though apparently constant, seasonal effects for all of the other series as well.

Table 4.3 gives the final estimates of the β_k parameters for the variable "employment rate" while Table 4.4 gives the same for "unemployment rate". We also did χ^2 tests of significance of two hypotheses of interest for these parameters; first we tested whether all of the β_k, k=2,...,6 were equal to each other, and if that test was non-significant we then tested whether the common value could be equal to zero. Tables 4.3 and 4.4 also include the p-values obtained from these tests. The hypothesis of a common value for the β_k parameters for k=2,...,6 was of interest because the level of non-response is higher for the first month in sample and the reasons for non-response might also be quite different (Kennedy, Drew and Lorenz, 1994).

Table 4.3 Summary of β_k estimates for the employment rate ($\times 10000$)

	Atl.	Que.	Ont.	Pra.	B.C.
β_1	-28	-24	-5	-21	-8
β_2	-14	-16	2	-5	-15
β_3	9	5	5	0	-12
β_4	5	8	-1	3	12
β_5	11	15	-3	7	7
β_6	16	13	2	16	16
p_1	0.005	0.088	0.684	0.459	0.258
p_2	n.a.	0.259	0.273	0.009	0.582

p_1 = p-value for test of $\beta_k = \beta$, k=2,...,6
p_2 = p-value for test of common $\beta = 0$

For the variable "employment rate" we found significant month-in-sample effects for only two of the five regions: the Atlantic and the Prairie regions; nevertheless, with the exception of Ontario, where all of the estimates are very close to zero, there is a distinct pattern in the estimated β parameters in that they tend to increase from negative to positive as the number of months in sample increases.

Table 4.4 Summary of β_k estimates for the unemployment rate ($\times 10000$)

	Atl.	Que.	Ont.	Pra.	B.C.
β_1	-19	-22	7	2	25
β_2	13	20	12	1	35
β_3	1	5	3	16	15
β_4	1	-2	0	6	-22
β_5	2	-11	-3	-8	-25
β_6	3	10	-19	-17	-27
p_1	0.841	0.188	0.120	0.018	0.011
p_2	0.004	0.000	0.970	n.a.	n.a.

p_1 = p-value for test of $\beta_k = \beta$, k=2,...,6
p_2 = p-value for test of common $\beta = 0$

For the variable "unemployment rate" we found significant month-in-sample effects in all regions except Ontario. Here the patterns in eastern and western Canada seem to be different. In both the Atlantic and Quebec regions, common positive values for β_k, $k = 2,\ldots,6$ are found, while in both the Prairies and B.C. the βs seem to show a trend from positive to negative.

4.2. Month-in-Sample Indices

Tables 4.5 and 4.6 show the month-in-sample indices, as defined in (3.1), for each of the five regions and for the whole of Canada, averaged over the 47 different occasions for which they could be calculated. In order to make more direct comparisons between the results obtained from the two approaches, we show in these Tables the differences of the month-in-sample indices from 1, multiplied by the 52-month averages of the estimated rates for each region. We also tested for each region and month-in-sample group whether the estimated average indices were significantly different from 1 by assuming that the 47 replicates were independent and applying a t-test.

In terms of the presence or absence of month-in-sample effects, the two approaches, state space modelling and month-in-sample indices, yield fairly consistent results.

From Table 4.5 we see that using the month-in-sample indices approach, significant month-in-sample effects were found for the variable "employment rate" for the Atlantic, Quebec and Prairie regions. From Table 4.3 the state space modelling approach yielded significant results for only the Atlantic and Prairie regions. However, it is worth noting from Table 4.3 that the test for a common month-in-sample effect for months 2 through 6 in Quebec yielded a p-value of 0.088 which is almost significant at the 5% level.

Table 4.5 Average month-in-sample indices for the employment rate: $(I_k - 1)\times\hat{r}\times 10000$

k	Atl.	Que.	Ont.	Pra.	B.C.	Can.
1	-29*	-29*	-5	-22*	-3	-16*
2	-15*	-16*	-2	-3	-13	-8*
3	8	5	3	0	-10	2
4	5	7	0	1	9	4
5	13*	17*	-2	8	5	7*
6	18*	16	6	17*	12	11*

* significant at the 5% level based on t-test

Table 4.6 Average month-in-sample indices for the unemployment rate: $(\hat{I}_k - 1)\times\hat{r}\times10000$

k	Atl.	Que.	Ont.	Pra.	B.C.	Can.
1	-19*	-21*	9	2	22	-1
2	18*	21*	17*	-1	37*	17*
3	5	7	4	16*	11	8*
4	0	-1	-2	10*	-18	-2
5	1	-13	-5	-9	-21*	-9*
6	-5	8	-23*	-19*	-30*	-13*

* significant at the 5% level based on t-test

From Table 4.6 we see that the month-in-sample indices gave significant results for all five regions for the variable "unemployment rate". From Table 4.4 the state space modelling approach gave significant results for all of the regions except Ontario, where the p-value for the test of a common month-in-sample effect for months 2 through 6 was 0.12.

Now comparing Tables 4.3 and 4.5, and Tables 4.4 and 4.6, we see that the two approaches yield very similar estimated month-in-sample effects for both the variables "employment rate" and "unemployment rate" within each region. Thus the two measures of month-in-sample effects validate each other.

The estimates of month-in-sample effects for the variable "employment rate" for the whole of Canada from Table 4.5 indicate a pattern of increasing from negative to positive as the number of months in sample increases. This pattern is also seen fairly consistently within each of the regions.

Considering now the estimated month-in-sample effects for the variable "unemployment rate" in Table 4.6, we see in all five regions a generally decreasing trend in the effects as k, the month-in-sample, increases regions, from 2 to 6; however, in the Atlantic and Quebec regions, the effect for $k=6$ is positive and much larger than that for $k=1$, while in the other regions the effect for $k=6$ is negative and much smaller than that for $k=1$. On reflection we may discern the same sort of pattern in Table 4.4, even though we concluded there from our tests of significance that common positive values of β_k, $k=2,\ldots,6$, are plausible for the Atlantic and Quebec regions. The estimates for the whole of Canada display a very smooth trend from positive to negative as k increases from 2 to 6.

5. DISCUSSION

As we have seen in Section 4, the two approaches based on state space modelling and month-in-sample indices yield generally similar patterns of estimated month-in-sample effects. This similarity of estimates from two different approaches suggests that both approaches are valid.

For the "employment rate" the largest estimated month-in-sample effects occur in Table 4.5 for the Atlantic and Quebec regions for the first month in sample, where the estimated effects are -0.0029. Assuming that the true value of "employment rate" lies somewhere in the range of the expected values of the estimates for month-in-sample groups, we may then conclude that the maximum possible bias in the estimate is 0.0029 and the true bias is probably much smaller. For the whole of Canada the maximum possible bias from Table 4.5 is 0.0016. A possible bias of this magnitude may be of some concern since it represents almost 40,000 persons at the national level; however, the bias would be roughly constant over time so that month to month comparisons of employment would not be greatly affected. Furthermore, in this paper we have studied the month-in-sample effects for sub-weighted estimates which make adjustments for household non-response within geographic areas. Recently the LFS has started to make non-response adjustments within month-in-sample groups. The actual estimates produced by the LFS are also calibrated to a variety of demographic characteristics such as population by age-sex groups within provinces. To the extent that the differences among estimates of "employment rate" and "unemployment rate" for different month-in-sample groups are explained by these factors, the possible bias in the published LFS estimates is that much smaller.

For the "unemployment rate" the largest estimated month-in-sample effect occurs in Table 4.6 for B.C. for the second month in sample, where the estimate is 0.0037. This suggests that the maximum possible bias in the estimate is 0.0037. At the Canada level the maximum possible bias from Table 4.6 would be 0.0017. Considering that changes in the unemployment rate of as little as 0.001 are generally newsworthy, a possible bias of this magnitude would be of some concern. However, the true bias is probably much smaller and would be fairly constant over time, and the bias of the published estimates, which are based on non-response adjustment within month-in-sample groups and are calibrated to various demographic totals, is likely much smaller again. Nevertheless, our results suggest that it would be worthwhile to estimate the month-in-sample indices for the estimators based on final weights to evaluate the possible extent of bias due to month-in-sample effects for the LFS estimates.

ACKNOWLEDGEMENT

We are grateful to Brian Kennedy for assistance in obtaining the data for this study and to Jack Gambino for comments on an earlier draft.

REFERENCES

Bailar, B. (1975). "The effects of rotation group bias on estimates from panel surveys", *Journal of the American Statistical Association*, 70, 23-30.

Kennedy, B., Drew, J.D., and Lorenz, P. (1994). "The impact of nonresponse adjustment on rotation group bias in the Canadian Labour Force Survey", presented at the 5th International Symposium on Household Survey Nonresponse, Ottawa, Canada.

Mohl, C. (1991). Internal memo, Statistics Canada.

Pfeffermann, D., and Bleuer, S.R. (1993). "Robust joint modelling of labour force series of small areas", *Survey Methodology*, 19, 149-163.