WORKING PAPER TSRA-89-016E

TIME SERIES RESEARCH & ANALYSIS DIVISION
METHODOLOGY BRANCH

A STUDY ON ACHIEVING ADDITIVITY OF SEASONALLY ADJUSTED EMPLOYMENT SERIES

BY

MARIETTA MORRY

September 1989

Time Series Research
and Analysis Division

Statistics Canada

ABSTRACT

The seasonally adjusted Canada Total Employment series can be obtained by adding up the seasonally adjusted component series according to different types of disaggregation such as age-sex breakdown, job duration by sex breakdown, geographic breakdown etc. Due to the non-linearities inherent to the seasonal adjustment method these totals do not coincide. This paper discusses two alternatives for reducing the discrepancy among the totals, one by using the additive decomposition during the seasonal adjustment of each component series, the other by forcing the components to add up to a common total through raking.

RÉSUMÉ

La désaisonnalisation de la série d'emploi total au Canada peut se faire indirectement auquel cas l'ajustement résulte de l'addition de composantes désaisonnalisées. La composition peut se faire selon différents regroupements: l'âge et le sexe, la durée d'emploi selon le sexe, un découpage géographique, etc. Cependant la méthode de désaisonnalisation n'étant pas linéaire, les totaux désaisonnalisés de ces regroupements ne seront pas égaux. Cet article présente les mérites respectifs de deux façons de réduire ces écarts: 1) désaisonnaliser chacune des composantes à l'aide du modèle additif, ou encore 2) forcer l'égalité des totaux en utilisant la technique de l'ajustement proportionnel itératif.

1.0 Introduction

This investigation was prompted by the large discrepancy observed in June between the seasonally adjusted total employment figures derived from four different aggregations most notably from the two aggregations one according to age-sex breakdown and the other one according to full-time part-time employment by sex. According to the first aggregation employment in June rose by 57,000 while the corresponding rise according to the second aggregation was only 13,000. Discrepancies between the two sets of figures have been noted at previous time points too but their magnitude was generally smaller. This deviation in seasonally adjusted totals is a direct consequence of the non-linearities inherent to the seasonal adjustment procedure such as multiplicative decomposition, the identification of extremes and the variable trend cycle routine. The problem of aggregation discrepancies is especially apparent during months in which the seasonal pattern of the component series varies greatly, as in the month of June, when part-time employment is considerably below the yearly average while full-time employment and employment according to the four age-sex categories are at a relatively high level.

Ideally it would be desirable to keep the seasonally adjusted totals fairly close to avoid the release of confusing signals by our agency. In the analysis to follow only two sets of series are investigated because their totals tend to deviate the most, i.e. employment by age-sex breakdown and employment according to full-time and part-time by sex. Two possible alternatives are considered one designed to reduce the size of discrepancies by selecting additive decomposition, the other, on the other hand would completely eliminate the problem by raking the full-time part-time seasonally adjusted employment figures to the more reliable age-sex breakdown total.

2.1 Switching to Additive Decomposition

As mentioned previously one source of the discrepancy is the application of multiplicative decomposition to the majority of the component series. Multiplicative decomposition is used when the amplitude of the seasonal component is not independent of the trend-cycle but increases or decreases in accordance with the change in the trend-cycle movement. If additive decomposition is applied to a series with multiplicative structure, the seasonal adjustment will yield a rapidly moving additive seasonal component in an attempt to force the seasonal to follow the movement in the trend-cycle. Since X11ARIMA has a tendency to distort seasonal movement at the end of the series the resulting seasonally adjusted data will undergo large revisions that could be avoided if the more stable multiplicative seasonal factors were applied.

At the present only the group Employed Men 15 to 24 and Male Part-time Employment are adjusted additively. In order to see if the gap between the seasonally adjusted totals could be reduced we forced additive adjustment for all the series. The May to June movement according to the age-sex breakdown remained virtually unchanged at +61,000 but there was a notable upward revision in the movement of the total of full-time part-time employment to 30,000 thus narrowing the gap to 31,000 between the two totals. Let us see the consequences of effecting this considerable change in the full-time part-time seasonally adjusted total. The X11ARIMA program calculates two statistics to indicate the amount and type of seasonality present in a series. The first statistic is the F-value from a standard one-way analysis of variance test applied to the seasonal-irregular component to measure if there is significant <u>stable seasonality</u> present in relation to the irregular fluctuations in the series. The higher the F stable value the stronger (and thus better estimated) is the seasonal component. The second F-value is obtained from a two-way analysis of variance test on the same data and it measures the amount of year to year movement in the seasonal component. Higher significant F-values are indicative of more rapidly moving seasonality and consequently of larger revisions in the seasonally adjusted figures.

Table 1.1 gives the F-values for stable and moving seasonality obtained from the additive and multiplicative seasonal adjustment of the four full-time part-time series.

Table 1.1 Stable and Moving Seasonality Test Results for Full-time and Part-time Employment Series

Series	Additive Decomposition		Multiplicative Decomposition	
	Stable Seasonality F-Test	Moving Seasonality F-Test	Stable Seasonality F-Test	Moving Seasonality F-Test
Men Full-time	2938*	1.16	3047*	1.91
Men Part-time	2181*	3.85*	2071*	5.63*
Women Full-time	709*	6.59*	1116*	.92
Women Part-time	186*	5.70*	348*	1.58

* Significant at the 1% level

Clearly, additive adjustment introduces significant moving seasonality and diminished stable seasonality to both the women full-time and part-time employment series signalling potential trouble with revisions. The series men part-time employment gives better F-values with additive adjustment but it is already adjusted additively.

The Men full-time employment series shows no significant moving seasonality with neither additive nor multiplicative decomposition. In this case one could switch to additive decomposition without the fear of introducing higher revisions even if the stable seasonal component is diminished somewhat as indicated by a lower stable seasonality F-value.

To measure the impact of switching to additive adjustment for the two women employment series a revision analysis was carried out on these two series using first the multiplicative and then the additive decomposition. Table 1.2 gives the results of the revision analysis.

Table 1.2 Comparison of the Size of Revisions using Multiplicative and Additive Decomposition

Series	First Year Revisions		Third Year (final) Revisions	
	Multiplicative	Additive	Multiplicative	Additive
Women Full-time	0.10	0.11	0.09	0.19
Women Part-time	0.32	0.38	0.36	0.55

From Table 1.2 one can see that the size of the average revisions (expressed as the absolute percentage change between the first published and subsequently revised seasonally adjusted series) is smaller using multiplicative decomposition. The difference in revision size is not very pronounced between first year revisions but it is substantial in terms of the final revision, e.g. for full-time women it more than doubled. Since it is the size of the final revision that is most relevant when judging the reliability of the first published seasonally adjusted figures it is clear that the additive decomposition yields much less reliable seasonally adjusted figures than the multiplicative one.

Thus only the men full-time series can be switched to additive adjustment without loss of reliability. This measure would narrow the gap between the month-to-month movements in the two totals but only by 9,000.

2.2 Forcing the Two Seasonally Adjusted Totals to Agree by Using Raking

Historically, the seasonally adjusted employment series according to age-sex categories have always been more reliable than the full-time part-time employment series by sex as indicated by smoother seasonally adjusted figures and smaller revisions. Thus it is reasonable that the full-time part-time series should be modified to add up to the age-sex total. In the procedure shown here the totals were forced to agree by raking the seasonally adjusted male full-time part-time series to the sum of the two age groups of the male employment series. Similar raking was carried out for the women full-time part-time series. Raking, otherwise known as iterative proportionate fitting, coincides with the generalized least squares method where the weight of each component is given by the actual value of the series and there is only one

constraint regarding the total. In effect, since this is only a one-dimensional problem there is no need for iteration, i.e. each component is multiplied by the corresponding percentage deviation between the two totals. Raking was carried out at the sex breakdown subtotals to ensure agreement not only at the all employees total level but also at this level of disaggregation (for the same cost).

As a measure of the impact of the raking on the component series the following statistics were calculated.

<u>M</u>ean <u>P</u>ercentage <u>D</u>ifference

$$MPD = \left(\sum_t \frac{CI_t - CI_t^R}{CI_t} \times 100\right) / n$$

<u>M</u>ean <u>A</u>bsolute <u>P</u>ercentage <u>D</u>ifference

$$MAPD = \sum_t \left|\frac{CI_t - CI_t^R}{CI_t} \times 100\right| / n$$

<u>MAX</u> <u>A</u>bsolute <u>P</u>ercentage <u>D</u>ifference

$$MAXAPD = \max_t \left|\frac{CI_t - CI_t^R}{CI_t} \times 100\right|$$

<u>M</u>ean <u>D</u>ifference-in-<u>G</u>rowth rate

$$MDG = \sum_t \left(\frac{CI_t - CI_{t-1}}{CI_{t-1}} - \frac{CI_t^R - CI_{t-1}^R}{CI_{t-1}^R}\right) \times 100 / n-1$$

<u>M</u>ean <u>A</u>bsolute <u>D</u>ifference-in-<u>G</u>rowth rate

$$MADG = \sum_t \left|\frac{CI_t - CI_{t-1}}{CI_{t-1}} - \frac{CI_t^R - CI_{t-1}^R}{CI_{t-1}^R}\right| \times 100 / n-1$$

MAX Absolute Difference-in-Growth rate

$$MAXADG = \max_t \left| \frac{CI_t - CI_{t-1}}{CI_{t-1}} - \frac{CI_t^R - CI_{t-1}^R}{CI_{t-1}^R} \right| \times 100$$

Mean Absolute Difference-in-Yearly-Growth rate

MADYG = same as MADG with t-1 replaced by t-12.

MAX Absolute Difference-in-Yearly-Growth rate

MAXADYG = same as MAYADG with t-1 replaced by t-12.

where CI_t - seasonally adjusted value at time t

CI_t^R - raked seasonally adjusted value at time t.

Table 2 shows the compiled statistics for each of the four raked component series full-time employment by sex and part-time employment by sex.

Table 2. Selected Summary Statistics on the Raked Employment Series

	Men Full-time	Men Part-time	Women Full-time	Women Part-time
MPD	0.01	-0.01	0.02	0.03
MAPD	0.08	0.07	0.09	0.14
MAXAPD	0.28	0.35	0.33	0.54
MDG	-0.01	-0.02	-0.02	-0.04
MADG	0.12	0.13	0.19	0.23
MAXADG	0.31	0.34	0.58	0.62
MADYG	0.04	0.04	0.07	0.06
MAXADYG	0.14	0.13	0.28	0.27

According to Table 2 the two female employment series were affected considerably more by the raking procedure than the male employment series and within each sex the part-time series changed more than the full-time series after the raking. In terms of level the average absolute percentage change (MAPD) introduced through raking was basically negligible around 0.1% for all the series with a maximum of 0.5% for part-time women. The month-to-month percentage movements however, underwent substantial modifications as shown by the mean absolute difference-in-growth rate(MADG). The average modification ranging in absolute value from 0.12 to 0.23 in the four series is quite high considering that the usual month-to-month absolute change is around 0.3. The problem is even more pronounced in certain month where the movement can be changed by a maximum of 0.62 (MAXADG).

To illustrate the consequence of modifying the month-to-month movement in the seasonally adjusted series consider to women full-time series in 1989 January. The original seasonally adjusted growth of 0.7% from December to January was changed to 1.1% after the raking. Lets examine if this growth rate is justifiable given the movements in the raw series. The December to January movements in the unadjusted series for the years 1982 to 1988 were the following: -1.2%, -0.1%, -3.1%, -1.1%, -1.4%, -2.7%, -1.0%. This means that the average December to January movement (which can be regarded as a rough estimate of the seasonal movement and the average monthly trend) is around -1.5%. Since this series increases about 3.6% on average per year the average monthly increase in trend is about 0.3%. Subtracting this value from the average month-to-month movement -1.5% yields -1.8% as the estimate of the average seasonal movement. Seasonal adjustment removes this average movement, thus for the movement from 1988 December to 1989 January we should obtain approximately +0.8% (-1.0-(-1.8)=+0.8) as the seasonally adjusted increase. In effect X11ARIMA estimated 0.7% as the seasonally adjusted growth rate very close to our rough estimate. However, the 1.1% growth obtained after raking is definitely high given the history of the series.

Another important statistic MADYGR measures the size of distortions introduced into the year-over-year movement by the raking procedure. Here again the distortions were higher for the female employment series with an average of 0.07% for full-time women, however this relatively high change in the rate is negligible in the context of a 3.5% average year-to-year growth rate. It is questionable though if one can ignore the change in these movements in certain months.

Users of seasonally adjusted data quite often check the year-over-year movement in the raw series to obtain an idea of the seasonally adjusted yearly growth rate. This practice is quite legitimate since basically all the seasonal movement would be cancelled out in a year-to-year comparison and the year-over-year growth rate in the raw is expected to be very close to that in the seasonally adjusted series.

For example the same month year ago comparison in the full-time women series yielded a 4.59% growth rate in July 1986 in the raw data while the corresponding figure in the seasonally adjusted series was slightly higher but comparable at 4.62%. The question is whether this movement is retained after raking. In 1985 July, raking introduced a -0.31% change to the seasonally adjusted data but in 1986 the figures were changed by only -0.04% after raking. The overall result was that the July to July movement in the raked series increased to 4.92% from the original 4.59%. This type of inconsistency could justly mystify users of the data.

There are several other factors one must consider when judging the acceptability of the raking procedure. Since we are dealing with seasonally adjusted series it is crucial that the raked series does not display any residual seasonality. It is also important to ensure that the raked series remain at least as smooth as the original ones.

To rule out the existence of residual seasonality the four raked component series were processed through the X11ARIMA program once again. The built-in F-tests did not detect the presence of any residual seasonality. Consequently the raking procedure was deemed acceptable from the point of seasonality.

Concerning smoothness the quality of the total was actually improved by forcing it to agree with the smoother age-sex breakdown total. The question is whether the raked component series were also acceptable in terms of smoothness. The statistic used for measuring smoothness was the standard deviation of the first-differenced seasonally adjusted series. This statistic was calculated for all four series before and after raking. The values remained virtually the same in all four cases, i.e. none of the components underwent a deterioration in smoothness.

The above exercise demonstrates that on average the raked component series are as reliable as the original seasonally adjusted series, yet in some months serious problems could be encountered concerning distortions in month-to-month movements and also in year-over-year movements.

3.0 Conclusion

Two alternatives were considered to reduce the discrepancy between seasonally adjusted employment totals obtained through two separate aggregations.

The first method managed to decrease the discrepancy by eliminating one of the non-linearities of the X11ARIMA method through switching to additive decomposition. This procedure, however introduced significant moving seasonality among some of the components leading to larger revisions.

The second alternative completely eliminated the discrepancy through raking the less reliable full-time part-time employment series to the more reliable age-sex breakdown total. Although on average the quality of the seasonally adjusted component series did not suffer due to raking, the month-to-month movements and year-over-year movements at some time points changed sufficiently to give reason for serious concern.

As a compromise, one could switch the Men full-time series to additive decomposition, this would not bring about a deterioration in quality and yet it would have narrowed the gap by 9,000 between the two totals in June.