Statistics Canada
Statistique Canada

Methodology Branch

Time Series Research and Analysis Division

Direction de la méthodologie

Division de la recherche et de l'analyse des chroniques

Canadä

Working Paper TSRA-85-043E

Cahier de travail SCRA-85-043E

Time Series Research and Analysis Division
Methodology Branch
Statistics Canada

Division de la Recherche et de l'analyse
des série chronologiques
Direction de la Méthodologie
Statistique Canada

ANALYSIS OF REVISIONS IN THE SEASONAL ADJUSTMENT OF DATA USING X-11-ARIMA MODEL-BASED FILTERS

by

Guy Huot, Kim Chiu and John Higginson, Statistics Canada
Nazira Gait, University of Sao Paulo, Brazil.

This is a preliminary version. Do not quote without authors' permission.
Comments are welcome.

ANALYSIS OF REVISIONS IN THE SEASONAL ADJUSTMENT OF DATA USING X-11-ARIMA MODEL-BASED FILTERS

by

Guy Huot

Kim Chiu

John Higginson

Statistics Canada

and

Nazira Gait

University of Sao Paulo, Brazil

Preliminary

Not to be quoted without authors' permission.

Invited Paper for presentation to the Fifth International Symposium on Forecasting, Montreal, June 9 - 12.

ABSTRACT

Concurrent seasonally adjusted values are subject to revision when more data become available. This study attempts to analyse the total revision associated with the concurrent seasonal filter for the X-11-ARIMA seasonal adjustment method. The total revision is defined as the mean-squared difference between the frequency response functions of the central and concurrent filters at certain frequencies. Four ARIMA models are considered which are used in the construction of the filter weights. We determine total revision for different forecast horizons and different ARIMA parameter values. Then we evaluate for different forecast horizons the sensitivity of total revision to change in model parameter values.

KEYWORDS: X-11-ARIMA, revision, frequency response function, region of improvement.

1. Introduction

Most data collected by a statistical agency are in the form of time series, that is a new datum point is collected every month or quarter, etc. In Canada, many time series contain annual or seasonal variation, and analysts often wish to have this seasonality removed before they study the series. The majority of seasonal adjustment methods adopted by statistical agencies belong to the category of moving-average procedures. The most widely used of these procedures is the X-11-ARIMA program with or without ARIMA extrapolations. The latter is referred to in this paper as the basic X-11.

The X-11-ARIMA procedure enhances the basic Bureau of the Census X-11 program by forecasting the raw time series one year ahead using parsimonious ARIMA models and by applying to the latest observed datum point an asymmetric seasonal adjustment filter which is closer to the central filter, thus reducing the filter revision for that point.

In practice, the figures seasonally adjusted using asymmetric filters will change, or be revised, as time passes. This revision has two causes; the addition of new data, and the consequent use of different moving averages or filters. In this paper, we study the revision due to the second cause only. Section 2 contains a mathematical definition of the filter revision used in X-11-ARIMA.

Some estimated parameter values in ARIMA models produce small filter revisions and other values, large revisions, sometimes larger than those of the basic Statistics Canada X-11 procedure. In section 3, we identify for four ARIMA models the sets of parameter values for which X-11-ARIMA has smaller filter revision than the basic X-11 procedure. In section 4, we discuss the parameter values obtained from fitting the four ARIMA models to about 190 Canadian economic time series, to see whether or not the fitted models produce revisions that are smaller than those of the basic X-11 procedure.

For a given time series, a given ARIMA model, and given parameter values, X-11-ARIMA may be better or worse than the basic X-11, depending on the forecast lead time, or forecast horizon. Section 5 considers this question.

2. Measure of revision associated with concurrent and central filters

Geweke (1978) and Pierce (1980) have shown theoretically that under certain conditions, the X-11-ARIMA seasonal adjustment procedure (Dagum, 1975, 1980) estimates a concurrent preliminary value whose total revision is smaller than the revision obtained using X-11 alone. The X-11-ARIMA proceeds as follows: (1) a univariate ARIMA model (see Box and Jenkins, 1970) is fitted to the series to be seasonally adjusted; (2) this series is extrapolated one year forward; and (3) provided the extrapolations are acceptable, the X-11 method is then applied to the extended series. That is, the current data are seasonally adjusted using weights that are closer to the central weights of the X-11 procedure on the extended series, rather than end weights.

Using end weights, that is, using a one-sided or asymmetric filter, as X-11 alone does, means that the seasonal adjustment of the last available figure, that is the concurrent figure depends on its past only. However, both the past and the future of the series being adjusted contain relevant seasonal information about the concurrent figure. Thus the ARIMA extension of the series allows for the use of past and acceptable forecasted future values when adjusting the concurrent figure.

The ARIMA model, by taking into account the particular autocovariance and variance structure of the series in making the forecasts, rapidly responds to its changing trend-cycle and seasonal variation. Combining the ARIMA extrapolation filter with the X-11 asymmetric filter for the seasonal adjustment of the concurrent point still results in an asymmetric filter, but with varying weights (see Dagum, 1983). This latter filter improves on the basic X-11 filter by reducing the mean squared error of the revisions. This holds, as shown in Pierce (1980), even for series that are non-stationary in their original form but stationary after a suitable differencing transformation.

The X-11 asymmetric filter for the seasonal adjustment of the concurrent point considered here assumes that the standard additive decomposition model is employed and that there is no replacement of extreme values.

Let x_j be the raw datum for time j and y_j be the analogous estimate of the seasonally adjusted number. The estimate of the seasonally adjusted value of the concurrent point x_j will be revised until the additional information is sufficiently distant from time j to be no longer relevant, in which case y_j is called central and its seasonal estimate is final. The successive revisions for y_j thus reflect both the quality of the new information and the differences in the weights applied to the series when subsequent relevant data $x_{j+\ell}$ are available. For our purposes, the data $x_{j+\ell}$ are acceptable ARIMA forecasted future values where $\ell \leq 42$.

The linear filters and revisions considered here are analysed in the frequency domain. Let the sequence of weights $\{h_k(c)\}_{-m}^{m}$ describe the filter used to seasonally adjust a point in a time series when it is far enough from the ends of the series to be central. $\{h_k(c)\}$ is then symmetric and its seasonal estimate is final. This symmetric filter is time-invariant. It is almost identical for X-11 and X-11-ARIMA since both methods vary mainly with proximity to the ends of the series. Let $\{h_k(c)\}_{-q}^{0}$ and $\{h_k(e)\}_{-q}^{0}$ be the weights of the X-11 and X-11-ARIMA concurrent point seasonal adjustment filters respectively. The latter will depend on the ARIMA model used, the parameter values, and the forecast horizon ℓ. Both filters produce a preliminary seasonal estimate. The properties of these filters have been studied in the frequency domain by Wallis (1974), Laroque (1977) and Dagum (1983) among other authors.

Any linear filter $\{h_k\}$ can be described in an equivalent way by means of its associated frequency response function $H(\omega)$:

$$H(\omega) = \sum_k h_k \exp(-i\omega k); \qquad (0 \leq \omega \leq \pi) \tag{2.1}$$

$$= G(\omega) \exp(i\Phi(\omega)) \tag{2.2}$$

where ω is the frequency,

$$G(\omega) = || \ H(\omega) \ || \tag{2.3}$$

the right hand side of the equation is the modulus of $H(\omega)$, and

$$\Phi(\omega) = \arctan \left[\frac{\mathrm{Im}(H(\omega))}{\mathrm{Re}(H(\omega))}\right] \tag{2.4}$$

The gain $G(\omega)$ and phase shift $\Phi(\omega)$ functions are particularly useful in the discussion of the properties of filters.

Given any periodic input of frequency ω radian per cycle to a filter $\{h_k\}$, the output will be amplified by a factor $G(\omega)$ and shifted by an angle $\Phi(\omega)$. For instance, while the symmetric filter produces a zero or $\pm\pi$ phase shift, the concurrent point seasonal adjustment filter can produces any shift between $\pm\pi$.

The revision measure, $R(\omega)$, is defined as

$$R(\omega) = H(\omega;e) - H(\omega;c), \tag{2.5}$$

that is, the difference between the frequency response functions of the concurrent point seasonal adjustment filters $H(\omega;e)$ and the symmetric filter $H(\omega;c)$. $R(\omega)$ takes into account both the gain and phase shift functions. The total revision measure, MSR, introduced by Dagum (1982), is defined as the mean square of the modulus of $R(\omega)$.

$$MSR = (1/\pi \int_0^\pi || R(\omega) ||^2 d\omega). \tag{2.6}$$

In the next sections, we will use the square root of MSR.

3. Stationarity and invertibility regions where the total revisions of X-11-ARIMA filters are smaller than those of X-11

We turn our attention, in this section, to the identification of the sets of parameter values of given ARIMA models for which the X-11-ARIMA current point seasonal adjustment filters improve on the X-11. The analysis of the X-11-ARIMA performance is undertaken for four models, namely

1. $(0,1,1)(0,1,1)_s$
2. $(0,1,2)(0,1,1)_s$
3. $(0,2,2)(0,1,1)_s$
4. $(2,1,0)(0,1,1)_s$

where s is 12 if the series is monthly and 4 if it is quarterly. The selection of the models resulted from a large empirical study based of Canadian macro-economic time series (Chiu, Higginson and Huot, 1984). This selection was obtained by taking into account the goodness-of-fit of the models, their forecasting performance and other criteria such as stability and invertibility conditions.

The four models will be incorporated in a new version of the X-11-ARIMA computer package. Only models 1, 3, and one other model are now available.

Figures 1,2 and 3 depict the admissible region for moving average (MA) models 1, 2 and 3 to be invertible. Figure 4 shows the stability region for model 4 that is autoregressive (AR). Stability and invertibility are dual concepts. As one can see in figures 2, 3 and 4, the stability region for an AR process is the same as the invertibility region for the analogous MA process.

FIGURE 1:
BOUNDARY AND ESTIMATED PARAMETER VALUES OF ACTUAL SERIES:
ARIMA MODEL $(0,1,1)(0,1,1)_s$

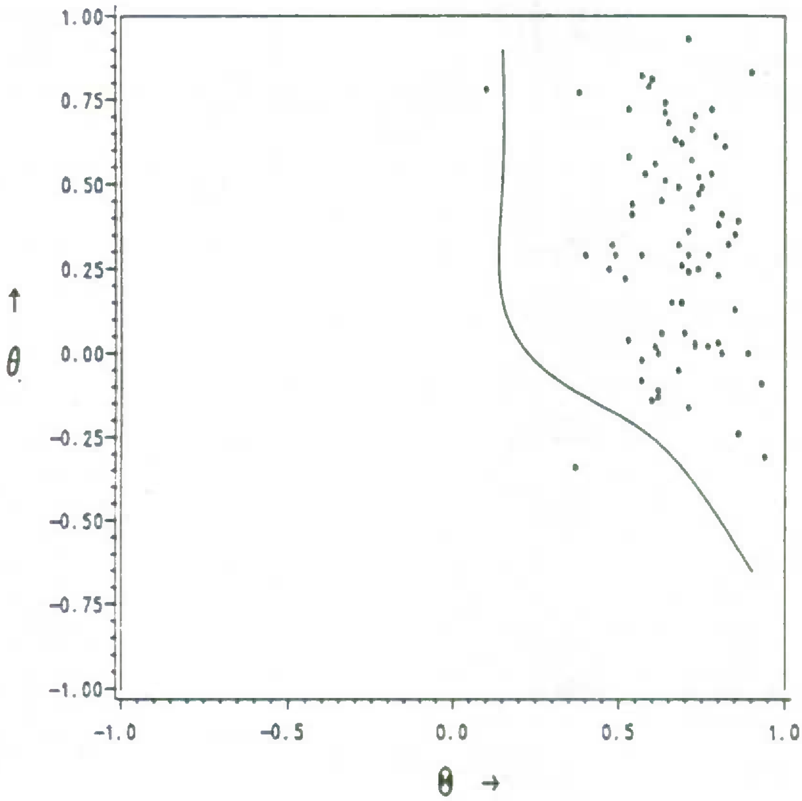

FIGURE 2:
BOUNDARY AND ESTIMATED PARAMETER VALUES OF ACTUAL SERIES:
ARIMA MODEL $(0,1,2)(0,1,1)_s$

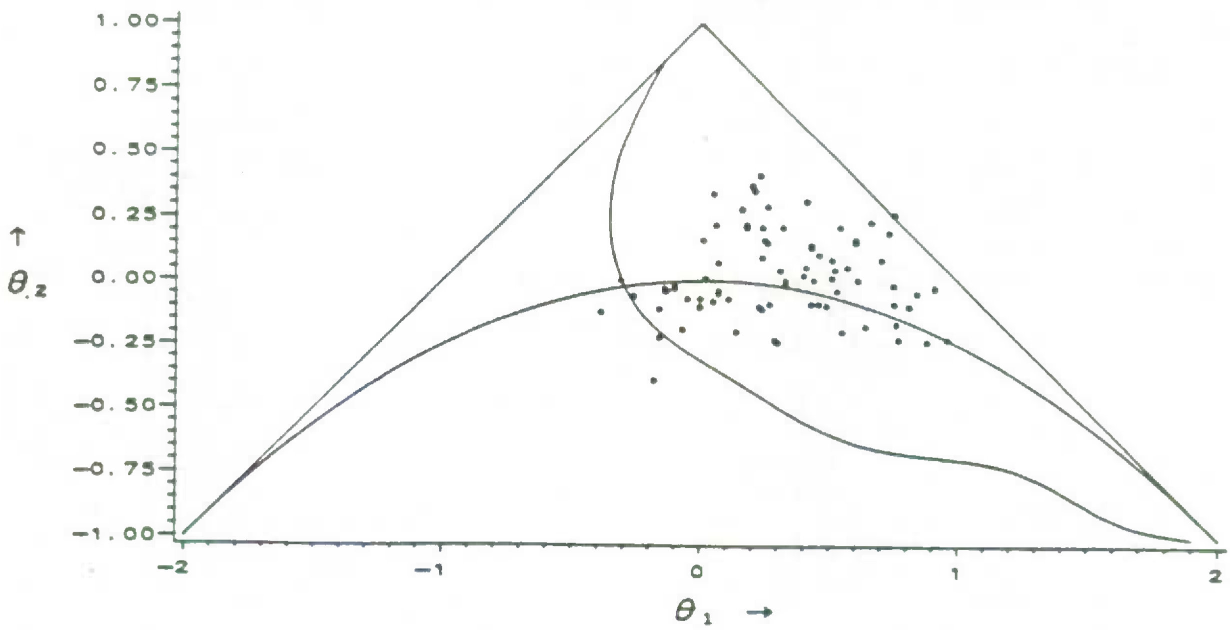

FIGURE 3:
BOUNDARY AND ESTIMATED PARAMETER VALUES OF ACTUAL SERIES:
ARIMA MODEL $(0,2,2)(0,1,1)_s$

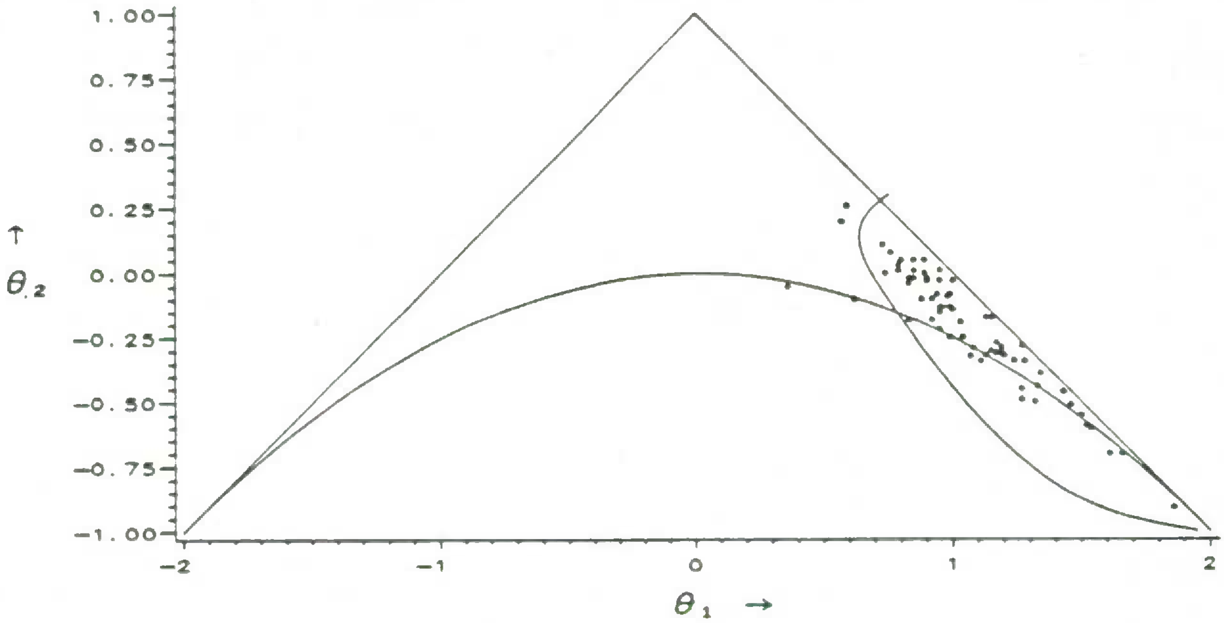

FIGURE 4:
BOUNDARY AND ESTIMATED PARAMETER VALUES OF ACTUAL SERIES:
ARIMA MODEL $(2,1,0)(0,1,1)_s$

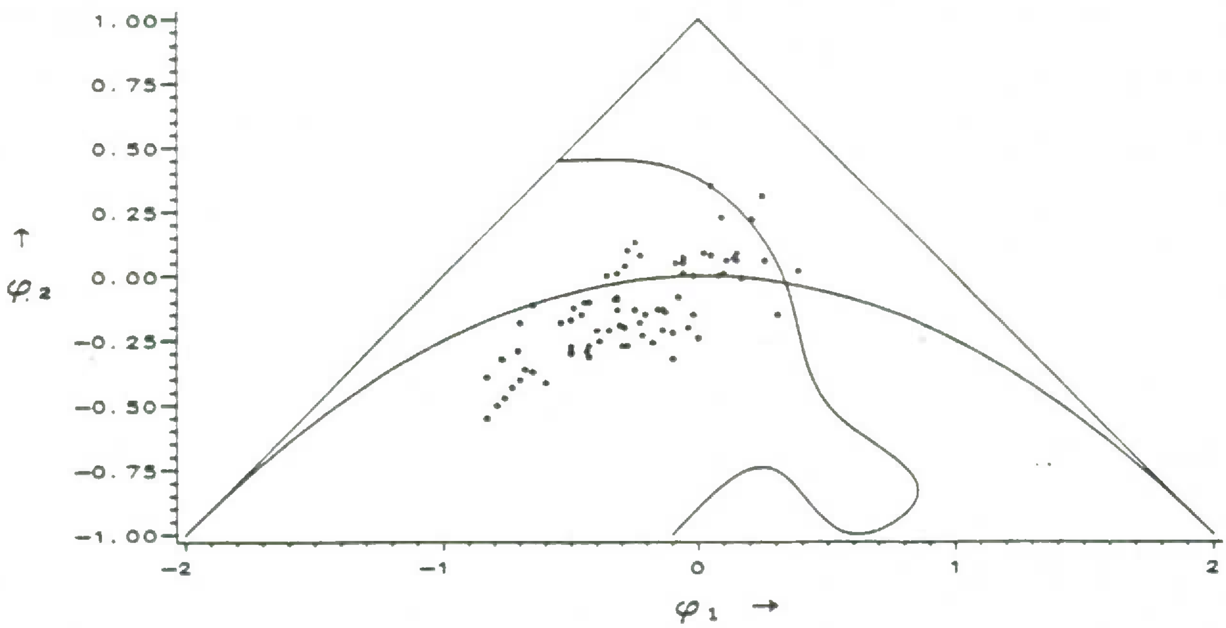

The four figures contain other pieces of information. We are here mainly concerned with the boundaries that define the set of ARIMA parameter values for which the X-11-ARIMA current point seasonal adjustment filters improve on the basic X-11 filter. The estimated parameter values for the sample of Canadian macro-economic time series, represented by the dots, and the parabolic boundaries in figures 2, 3 and 4 will be discussed in the next section.

The boundaries plotted on the invertibility and stability regions for the θ's and ϕ's are estimated by comparing the modulus of revision, MSR, for the basic X-11 procedure and for X-11-ARIMA. For each of the four models, the values of MSR were calculated for a grid of values of (θ, Θ) for model 1, (θ_1, θ_2) for models 2 and 3, and (ϕ_1, ϕ_2) for model 4. X-11-ARIMA improves on X-11 whenever its MSR is smaller than 0.359, which is the value of the modulus of revision for X-11 alone. The boundary lines indicate those pairs of parameter values having a MSR smaller than 0.359.

The structure of the series that fall in the region of improvement for X-11-ARIMA is generally the same as that implicit in the X-11 central weights. The ARIMA models force the X-11-ARIMA concurrent filters to converge faster to these weights. Series that fall outside the region are likely to have a different structure.

A complete family of boundaries exists for each of the four figures. The boundaries vary according to the forecast horizon ℓ, and in figures 2, 3 and 4, according to the value taken by Θ. In order to simplify the graphic presentation,

ℓ was set to 12 in all four figures and Θ was set to 0.6 in the last three; however a family of boundaries has been estimated for each model for ℓ = 6, 12, 18, 24 and 42 months and for Θ = 0.8.

The region of improvement of X-11-ARIMA on X-11 for parameters of the first model is the upper right part of figure 1. The boundaries for ℓ = 12 and 18 months are vertical in their upper half and Θ is always greater than or equal to 0.2. For ℓ = 6 months, $\Theta \geqslant -0.6$, and for ℓ = 24 and 42 months, Θ is $\geqslant$ 0.3. We see that Θ increases with ℓ. In other words, the smaller Θ is (i.e. the more the seasonal pattern changes), the shorter the forecast horizon must be in order for X-11-ARIMA to improve on X-11. On the other hand, when ℓ increases θ decreases.

The region of improvement for parameters of the second model is in the upper right of the triangle. When ℓ increases from 6 to 12 months, the boundary of the region moves significantly to the left. As ℓ increases to 18, 24, and 42 months, the leftward movement of the boundary is smaller. Increasing Θ from 0.6 to 0.8 always moves the boundary to the left.

The third model has a very small region of improvement in the lower right of the triangle. We shall see why this is so in the next section. In general the boundary of the region behaves like that for the second model.

The last model differs from the preceding ones in that its region of improvement is in the left part of the triangle. Now the boundary moves to the right when ℓ or Θ increases. This opposite movement occurs because this last model represents an autoregressive process while the other models represent moving-average processes.

Each model has a region where the X-11-ARIMA end point seasonal factor filters improve on those of X-11. From the point of view of filters, any series fitted and forecasted by an ARIMA model whose parameter values fall in this region should be seasonally adjusted better by X-11-ARIMA than by X-11.

4. The estimated parameter values for actual series

The most important finding shown in figures 1 to 4 is that almost all the estimated parameter values in our sample of Canadian macro-economic time series fall within the region where X-11-ARIMA performs better than the basic X-11. This means that for each of these series there is a gain in seasonal factor revision due to the filter.

The sample comprised 23 quarterly series and 167 monthly seasonal time series chosen randomly from eleven sectors of the Canadian economy; national accounts; labour; prices; manufacturing; fuel, power and mining; construction; food and agriculture; domestic trade; external trade; transportation; and finance.

Among the 190 series to which each of the models were fitted, we graphed in figures 1 to 4 only those series that satisfied the two extrapolation criteria of the X-11-ARIMA. That is, an adjusted model is acceptable if its forecast error is less than or equal to 15% and if the χ^2 probability for testing the null hypothesis of randomness of residuals is greater than or equal to 5%.

Table 1 shows, for each of the four models, the number of series that satisfy the two extrapolation criteria of X-11-ARIMA, the mean and standard deviation of the estimated parameters, their smallest and largest values and their range.

We see from figure 1 that only two of the 79 series that were well fitted and extrapolated by the first model, have parameters that fall outside the region. Models 2 and 3 each have 3 such series and model 4 has 4.

TABLE 1: Characteristics of the estimated parameter values for the four ARIMA models

Model	N	Variable	Mean	Standard Deviation	Min. Value	Max. Value	Range
$(0,1,1)(0,1,1)_s$	79	θ	0.33	0.31	-0.34	0.93	1.27
		Θ	0.67	0.14	0.10	0.94	0.84
$(0,1,2)(0,1,1)_s$	81	θ_1	0.31	0.32	-0.39	0.96	1.35
		θ_2	0.02	0.17	-0.39	0.41	0.80
		Θ	0.65	0.15	0.09	0.96	0.87
$(0,2,2)(0,1,1)_s$	67	θ_1	1.05	0.27	0.36	1.86	1.50
		θ_2	-0.20	0.23	-0.91	0.26	1.17
		Θ	0.60	0.15	0.02	0.92	0.90
$(2,1,0)(0,1,1)_s$	79	ϕ_1	-0.26	0.30	-0.83	0.39	1.22
		ϕ_2	-0.13	0.19	-0.55	0.35	0.90
		Θ	0.68	0.12	0.33	0.98	0.65

One may wonder, when looking at table 1, at the small number of series (at most 43%) that each model fitted and forecasted satisfactorily or at the average value of Θ for the different models (at most 0.68). The most probable explanation for the small value of Θ is that we used the conditional least-squares estimation method, instead of the unconditional method. As for the success rate of the ARIMA models, we should take into account the fact that these series ended in December 1983. The forecasts were obtained by fitting the models to the entire series in order to estimate the parameter values and calculate the forecasts for the last three years. An important reason for this high rate of failure of 57% is that the last three years were atypical years afflicted by a severe recession which has caused a rate of failure (between 34% and 41%) in forecast error. The rate of failure in χ^2 ranges between 22% and 29% (see: Chiu, Higginson and Huot, 1984).

The size of the improvement due to the filters varies within the region of improvement. That is, there is a point in this region where the filter revision is smallest. MSR is a strictly increasing function in all directions from the smallest point, throughout the entire region (of invertibility for models 1 to 3 and of stability for model 4). Outside the region of improvement, MSR increases much more rapidly than inside. We know the size of the filter improvement for each set of parameter values. Thus if several models fit a series fairly well, all other things being equal, we have enough information to select the model whose parameter values provide the smallest filter revision.

These results not only conform to but extend the conclusions given in several theoretical and empirical studies (see among others: Geweke, 1978; Pierce, 1980; Kenny and Durbin, 1982; Dagum, 1975, 1982.a and 1982.b; and Wallis, 1983). This is because the results show that as long as the extrapolations are reasonable (a forecast error $\leq$ 15%) and the fitting is acceptable ($\chi^2 \geq$ 5%), the extrapolations will generally reduce the concurrent filter revisions although the model chosen might not be the optimal one from the point of view of the mean squared-error of forecasts.

To illustrate the nature of the improvement due to the filters within the region of improvement, we will now analyse different sets of parameter values for each model using their frequency response function $H(\omega)$, their gain $G(\omega)$ and phase shift $\phi(\omega)$.

For our first model we chose two pairs of parameters inside the improvement region. The first, $\theta = 0.5$ and $\Theta = 0.9$, is the point where the filter revision is smallest. The second, $\theta = 0.4$ and $\Theta = 0.4$, is a point near the boundary.

Figure 5 shows three curves. Each represents graphically the distance between the frequency response functions of two filters. We are comparing the central X-11 filter with three different end filters in turn: the X-11 and those of X-11-ARIMA associated with the two chosen pairs of parameters. The lower the curve, the better the associated end filter is performing. The curve with the lowest average height is the one associated with (0.5, 0.9). The highest average height belongs to the curve for X-11.

From this figure we see that the X-11-ARIMA concurrent filter with parameter values $\theta = 0.5$ and $\Theta = 0.9$, on average, approximates best the frequency response function of the X-11 central filter. It improves on the second X-11-ARIMA concurrent filter because its gain function is closer to unity and it has less phase shift as shown in figures 6 and 7. It also improves, on average, on the X-11 concurrent filter especially at frequencies close to the seasonal frequencies k cycles per year, k = 1, ..., 6. The improvement occurs

FIGURE 5: DISTANCE BETWEEN CENTRAL X-11 FILTER AND 3 DIFFERENT END FILTERS.
THE END FILTERS ARE: X-11
: X-11 ARIMA (θ=.40 Θ=.40)
: X-11 ARIMA (θ=.50 Θ=.90)
ARIMA MODEL (0,1,1)(0,1,1)s - FORECAST HORIZON: 12

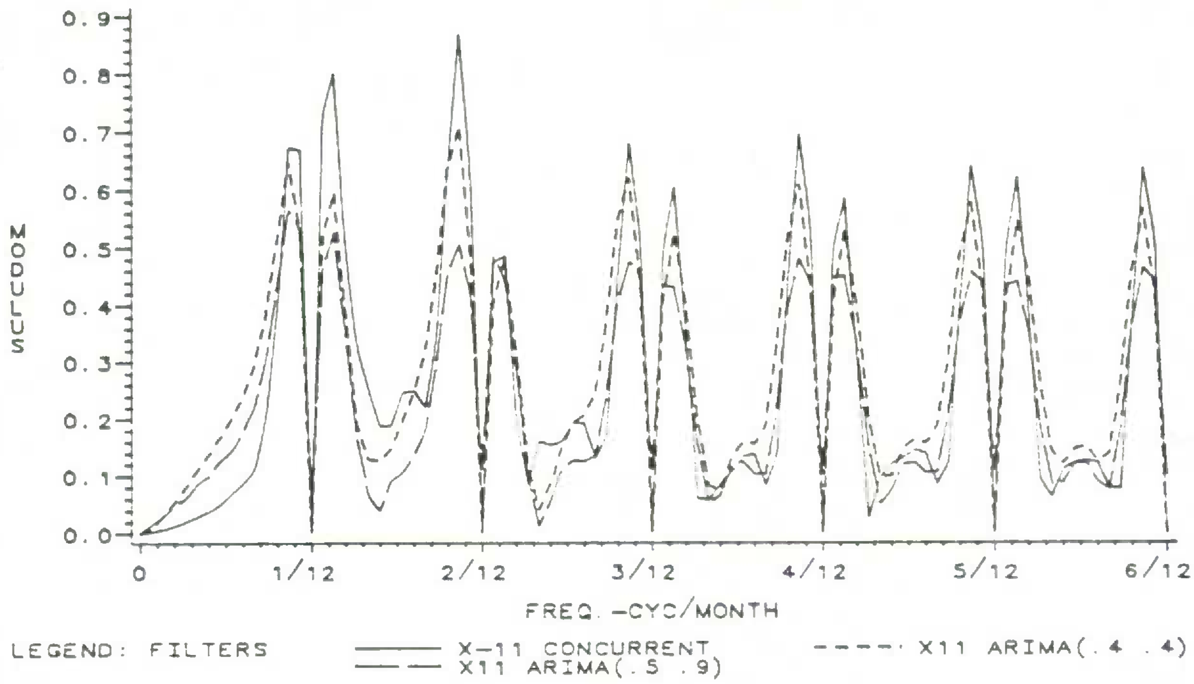

FIGURE 6: GAIN FUNCTIONS OF SEASONAL ADJUSTMENT FILTERS.
THE FILTERS ARE: X-11 CENTRAL
: X-11 CONCURRENT
: X-11 ARIMA (θ=.40 Θ=.40)
: X-11 ARIMA (θ=.50 Θ=.90)
ARIMA MODEL (0,1,1)(0,1,1)s - FORECAST HORIZON: 12

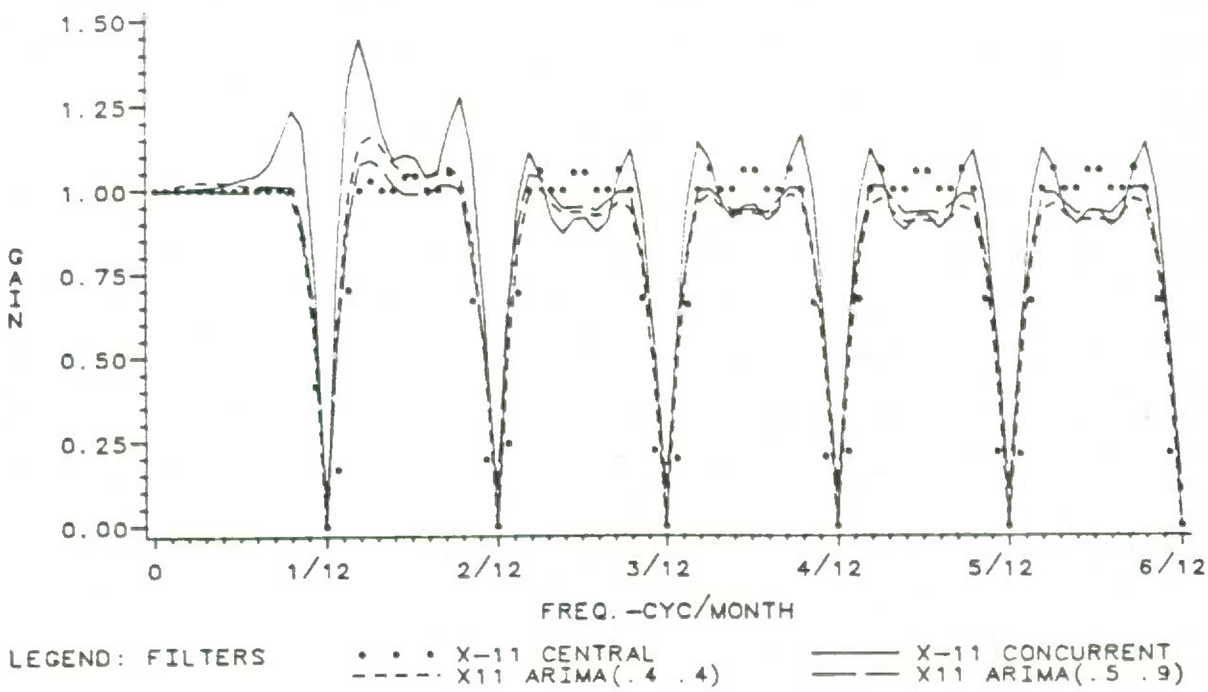

FIGURE 7: PHASE-SHIFT FUNCTIONS OF SEASONAL ADJUSTMENT FILTERS.

THE FILTERS ARE: X-11 CONCURRENT
: X-11 ARIMA (θ=.40 Θ=.40)
: X-11 ARIMA (θ=.50 Θ=.90)
ARIMA MODEL (0,1,1)(0,1,1)s - FORECAST HORIZON: 12

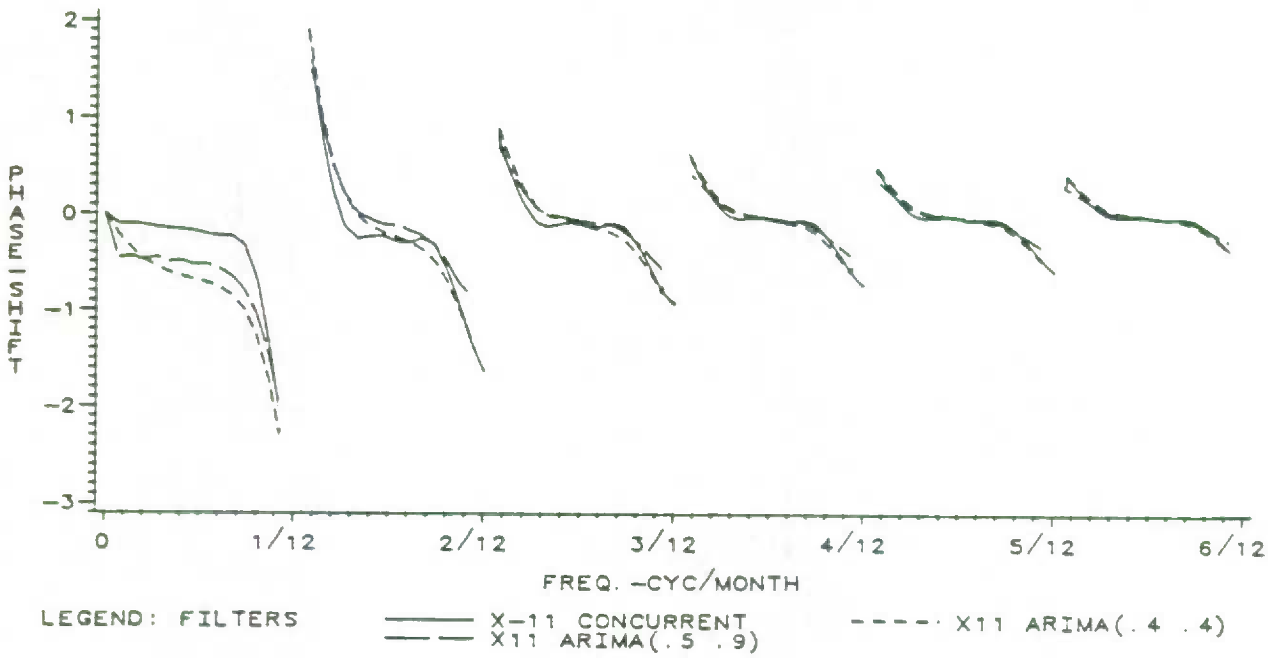

because the gain function for the ARIMA (0.5, 0.9) filter is much closer to unity at these frequencies. A gain function close to unity indicates that the contribution of these frequencies to the series is almost unaltered in power. However, both ARIMA filters have a poorer performance than the X-11 concurrent filter in the business-cycle range. This is due to a large phase shift.

The border of the improvement region for all four models occurs for a revision of 0.359. The minimum revision possible for model 1 in figure 1 is 0.285. Most of the series modelled in this figure have a revision between 0.300 and 0.330.

The roots of the characteristic equation of a second-order AR or MA process are real if ϕ_2 or θ_2 lie above the parabola shown in figures 2, 3 and 4, and complex if they lie below. We see from figure 2 that about one third of the series fitted by model 2 had complex roots.

Model 2 is characterized by a large region of improvement. The smallest revision, located in the upper right part of the region of improvement, is obtained with parameter values $\theta_1 = 0.4$ and $\theta_2 = 0.6$. The gain function associated with this filter is significantly lower than unity (viz 0.90). This performance in gain is more than compensated for in superior performance in phase shift and it is the filter that allows best for moving seasonality. Inside the region of improvement, except in its upper right part, the gain functions of the filters associated with most of the series are close to unity except of course at seasonal frequencies. Outside the region of improvement where few series fell, the gain of the associated filter is close to unity at business-cycle frequencies but rises significantly above unity at high non-seasonal frequencies. The minimum revision possible for an invertible model in figure 2 is 0.316. Most of the series shown in this figure had revision less than 0.330.

The region of improvement for model 3 is small and located to the lower right of the triangular invertibility region in figure 3. The equation of the right side of the triangle is $\theta_2 = 1 - \theta_1$. The cloud of points in figure 3 shows that the estimated value of θ_2 is usually close to the estimated value of $1 - \theta_1$. The minimum revision possible for the third model is 0.321 for $\theta_1 = 1.8$ and $\theta_2 = -0.8$. Most of the series modelled have a revision lower than 0.345.

There are two important differences between models 2 and 3. The gain function associated with the minimum revision filter for model 3 is equal to unity at business-cycle frequencies, which was not the case for model 2. The minimum revision filter for model 2 allows somewhat better than that of model 3 for moving seasonality. These two models have similar gain functions in other respects.

The minimum revision possible for model 4 is 0.324 for $\phi_1 = -0.3$ and $\phi_2 = -0.3$. Most series modelled in figure 4 have revision lower than 0.350. Unlike the other three models, the minimum revision point for this model is located almost in the middle of the region of improvement. Its gain function is closer to unity than that of models 2 and 3, and it allows for moving seasonality slightly less than these two models. It is not even the filter that allows best for moving seasonality in figure 4. Filters located to the left of $\phi_1 = -0.3$ and $\phi_2 = -0.3$ perform better from that point of view.

So far the analysis has focussed on the region of improvement. However, outside this region, is the X-11-ARIMA worse than X-11? In order to answer the question, we may reasonably assume that the few series that fall outside the region of improvement are likely to have a structure that is different from that implicit in the X-11 central weights. Cleveland (1972) and Cleveland and Tiao (1976) have shown that there exists a stochastic model, very close to the $(0,1,1)(0,1,1)_s$ model with parameter values $\theta = 0.4$ and $\Theta = 0.6$, that approximates well the central weights of the X-11 procedure. These parameter values lie inside the region of improvement for the first model. When we depart significantly from $\theta = 0.4$ and $\Theta = 0.6$, the true structure of series is not likely to be the same as that implicit in the X-11 central weights. Thus it is not desirable to have concurrent filters that converge to the X-11 central filter.

Let us assume that the ARIMA model fitted to such series gives a good representation of the behaviour of the series that lie outside the region of improvement. Then the concurrent filter of X-11-ARIMA, which is mainly influenced by the extrapolation filter of the ARIMA model, reflects the current structure of the series. In this case, the total revision for the concurrent filter will be larger than that of X-11. It does not mean that the concurrent estimate of X-11-ARIMA for these few series is not better than the one obtained from X-11 alone. In such a situation, one should revise the concurrent estimate at most one year later in order to avoid forcing the convergence toward the inappropriate X-11 central filter.

5. Revisions for different forecast horizons

We will now investigate the effect of different forecast horizons on the total revision, MSR, for the concurrent X-11-ARIMA filters. The forecast horizon ℓ denotes the lead time of the forecast. We will first compare the MSR associated with different forecast horizons $\ell = 1, \ldots, 24$ and 42 months for different pairs of parameter values located in the region of improvement of the first model. Then for each model we will study the forecast horizon for the concurrent seasonal filter of X-11-ARIMA associated with the smallest MSR value.

Figure 8 shows five curves, each corresponding to a point in the improvement region for model 1. The set of parameter values $\theta = 0.5$ and $\Theta = 0.9$ which is the point where the filter revision is smallest has the lowest curve for every lead time. We see furthermore from this figure that the optimal forecast horizon for revisions depends strongly on the location of the parameter values in the region of improvement. For $\theta = 0.5$ and $\Theta = 0.9$ the best lead time is 24 months. However, the longer the forecast lead time, the greater the variance of the forecast error and the riskier it is to forecast. In practice the lead time for the X-11-ARIMA is 12 months. It appears from figure 8 that there are other cases such as parameter values $\theta = 0.8$ and $\Theta = 0.4$, and $\theta = 0.4$ and $\Theta = 0.4$, where the lead time should not be extended beyond 11 months.

Twelve months is a very sensitive lead time. It can be associated in figure 8 with the largest decrease in MSR (0.5, 0.9) or the largest increase in MSR (0.8, 0.4). The lead time of 24 months bahaves in a similar manner.

The end point x_j is treated as central by X-11-ARIMA only when there are acceptable ARIMA forecasted values for 42 months ahead. Then x_j is seasonally adjusted using a symmetric filter. Figure 8 shows that there is not much gain to be expected by extending the forecast lead time from 24 to 42 months.

FIGURE 8: REVISIONS FOR DIFFERENT PARAMETER VALUES FOR ARIMA MODEL (0,1,1)(0,1,1)s

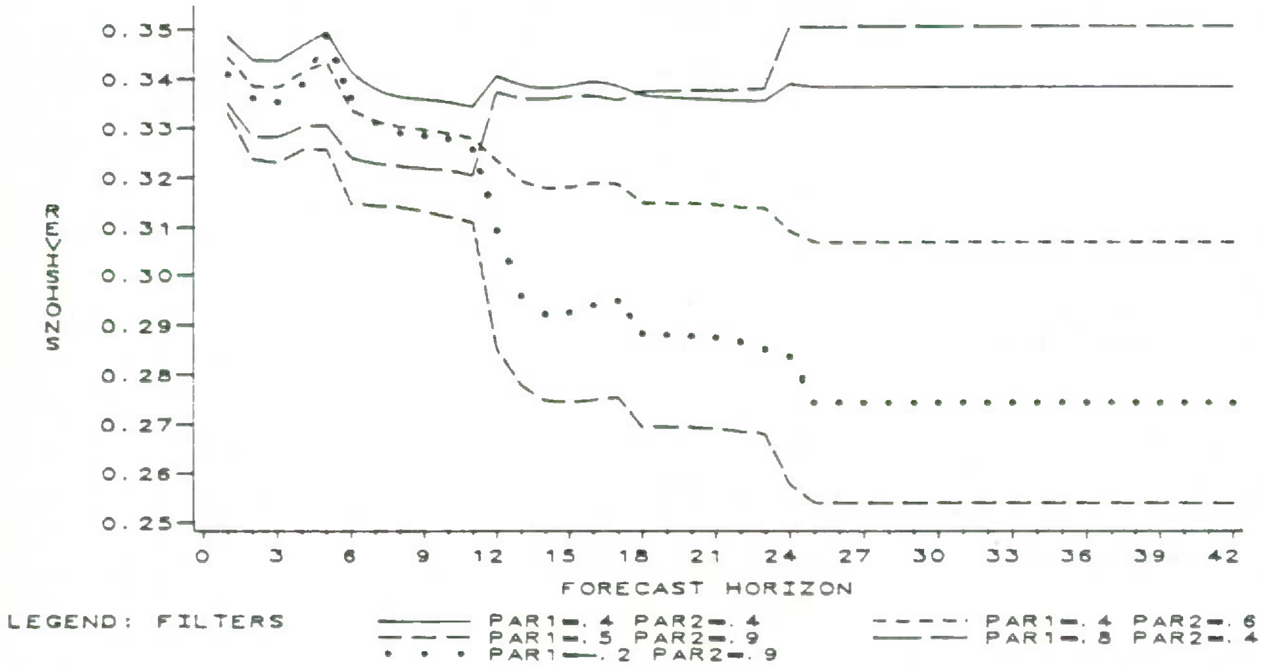

Table 2 shows, for each of the four models, and for lead times ℓ = 6, 12, 18, 24 and 42 months, the MSR values where the filter revision is smallest inside the regions of improvement. The first model gives the smallest revision. We see that the closer the seasonal parameter Θ is to 1.0, the smaller the revision. Such a value of Θ corresponds to stable seasonality, which is easy to forecast using ARIMA models.

Table 2: MSR values associated with the smallest filter revision inside the region of improvement for each ARIMA model

ℓ	$(0,1,1)(0,1,1)_s$ (.5, .9)	$(0,1,2)(0,1,1)_s$ (.4, .6)		$(0,2,2)(0,1,1)_s$ (1.8, -.8)		$(2,1,0)(0,1,1)_s$ (-.3, -.3)	
		$\Theta = .6$	$\Theta = .8$	$\Theta = .6$	$\Theta = .8$	$\Theta = .6$	$\Theta = .8$
6	.315	.321	.315	.320	.315	.335	.324
12	.285	.316	.298	.320	.301	.324	.301
18	.270	.311	.286	.315	.288	.316	.288
24	.258	.312	.278	.315	.280	.310	.277
42	.254	.307	.270	.311	.272	.308	.274

The other points in the region of improvement of models 2, 3 and 4 have a behaviour similar to that shown in figure 8. There exist sets of parameter values outside the region of improvement for which MSR increases as ℓ increases.

6. Conclusion

In this paper, we have shown that there exist sets of parameter values of given ARIMA models for which the X-11-ARIMA concurrent point seasonal adjustment filter improves on the basic X-11. These sets, referred to as the regions of improvement for X-11-ARIMA, are defined by boundaries plotted on the invertibility and stability regions for the θ's and ϕ's.

- From the point of view of filters, any series fitted and forecasted by an ARIMA model whose parameter values fall in the region of improvement will generally be seasonally adjusted better by X-11-ARIMA than by the basic X-11.
- An important finding is that almost all the estimated parameter values in our sample of Canadian macro-economic time series fall in the regions of improvement.
- The size of the improvement due to the filters varies within the region of improvement. There is a point in this region where the filter revision is smallest. The revision is a strictly increasing function in all directions from that point.
- We know the size of the filter improvement for each set of parameter values. Thus if several models fit a series fairly well, all other things being equal, we have enough information to select the model whose parameter values provide the smallest filter revision.
- As long as the extrapolations are reasonable (a forecast error $\leq$ 15%) and the fitting is acceptable ($\chi^2 \geq$ 5%), the extrapolations will generally reduce the concurrent filter revisions although the model chosen might not be the

optimal one from the point of view of the mean squared-error of forecasts.

- The optimal forecast horizon for revisions depends strongly on the parameter values in the region of improvement. There are cases where the best lead time is 24 months and other cases where it should not be extended beyond 11 months.
- We can select the forecast horizon according to the estimated parameter values.
- Twelve and 24 months are very sensitive lead time.
- From the point of view of filters, there is not much gain to be expected by extending the forecast lead time from 24 to 42 months.

REFERENCES

Box, G.E.P. and Jenkins, G.M. (1970): "Time Series Analysis Forecasting and Control", San Francisco; Holden Day.

Chiu, K., J. Higginson and G. Huot (1984), "Performance of ARIMA Models in Time Series", American Statistical Association, Proceedings of the Business and Economic Statistics Section, 425 - 430.

Cleveland, W.P. (1972): "Analysis and Forecasting of Seasonal Time Series, Ph.D. dissertation (Department of Statistics, University of Wisconsin, Madison, WI).

Cleveland, W.P. and G.C. Tiao (1976): "Decomposition of Seasonal Time Series: A Model for the Census X-11 Program", Journal of American Statistical Association, 71, 581-587.

Dagum, E.B. (1975): "Seasonal Factor Forecasts from ARIMA Models" Proceedings of the International Statistical Institute, 40th Session, 3, Warsaw, 206-219.

Dagum, E.B. (1980): The X-11-ARIMA Seasonal Adjustment Method, Ottawa: Statistics Canada, Catalogue No. 12-564E.

Dagum, E.B. (1982): "Revisions of Time Varying Seasonal Filters", Journal of Forecasting, Vol. 1, 173-187.

Dagum, E.B. (1983): "Spectral Properties of the Concurrent and Forecasting Seasonal Linear Filters of the X-11-ARIMA Method", Canadian J. of Stat., Vol. 11, No. 1, 73-90.

Geweke, J. (1978): "The Revision of Seasonally Adjusted Time Series", SSRI Report No. 7803, Dept. of Economics, University of Wisconsin.

Kenny, P.B., and Durbin, J. (1982), "Local Trend Estimations and Seasonal Adjustment of Economic and Social Time Series", Journal of the Royal Statistical Society, Ser. A, 145, part 1, 1-41.

Laroque, G. (1977): "Analyse d'une méthode de désaisonnalisation: le programme X-11 du U.S. Bureau of Census, version Trimestrielle", Annales de l'I.N.S.E.E., 28, 105-127.

Pierce, D. (1980): "Data Revision with Moving Average Seasonal Adjustment Procedures", Journal of Econometrics, 14, 95-114.

Wallis, K.F. (1974): "Seasonal Adjustment and Relations Between Variables", Journal of American Statistical Association, 69, 18-31.

Wallis, K.F. (1983): "Models for X-11 and X-11-FORECAST Procedures for Preliminary and Revised Seasonal Adjustment", in Applied Time Series Analysis of Economic Data, ed. A. Zellner, Washington, D.C.: U.S. Bureau of the Census, 3-11.

c.3