Issues in Trend-Cycle Estimates for Official Statistics

Begoña Martín[1], Peter Meszaros[2], Duncan Elliott[1], Craig McLaren[1]
[1]Office for National Statistics, Cardiff Road, Newport, United Kingdom, NP10 8XG
[2]Statistics Canada, 100 Tunney's Pasture Driveway, Ottawa, Ontario, K1A 0T6

Abstract
The current economic climate emphasizes the importance of identifying turning points in time series quickly and accurately. Trends are often used to describe underlying movements in series, so need to perform well in periods of changes to avoid misleading users. We examine the properties of short-term trend estimation methods with regards to accuracy, bias and revisions. We evaluate the performance of various filters – including the recent Cascade Linear Filter (Dagum and Luati, 2009) – and discuss the underlying motivation behind these measures.

Key Words: Trend estimation, Henderson filter, ARIMA, false turning points

1. Introduction

In many National Statistical Institutes (NSIs) it is standard practice to publish key economic time series in seasonally adjusted form. This is done to help users interpret series, particularly over the short-term. Seasonal adjustment first estimates and then removes the systematic calendar related components from the time series. However, the seasonally adjusted estimates still contain the trend and the irregular components. The presence of these irregular movements often makes interpretation difficult. Therefore, some NSIs also publish trend estimates, presenting users with their best estimates of the underlying behaviour of the series.

The practice of using trends and the methods for estimating trends varies considerably between the NSIs around the world. In the United Kingdom Office for National Statistics (ONS), the use of trends is very limited because there is a major concern around the reliability of the trend, particularly at the current end of the series, where trend estimates are very likely to be revised. Statistics Canada, which publishes trends along with their seasonally adjusted series, use alternative methods to the standard 13-term Henderson moving average from X-12-ARIMA (Findley et al., 1998) to try and address the problem of large revisions. The Australian Bureau of Statistics (ABS) regularly produces and publishes trend estimates for all of their main outputs (ABS, 2003).

Ideally, trend estimation methods should: accurately and quickly identify turning points, especially at the end of the series; minimize the number of false turning points; and have minimal and unbiased revisions. However, no trend estimation method will fulfil all these criteria at once. For example, Gray and Thomson (1996) give a general framework for deriving trend estimates and demonstrate that there is a certain trade-off between the desired properties. Research by the ONS (Compton, 2000) concluded that quick detection of turning points and minimisation of the number of false turning points were considered to be the most important characteristics of a trend for most NSIs, while the trend being unaffected by outliers was considered to be the least important characteristic.

In this paper, we evaluate the performance of three non-parametric trend predictors on a set of real time series; the Retail Sales Index from the UK and Canada. The Non Linear Dagum Filter (NLDF) developed by Dagum (1996), the Asymmetric Linear Filter (ALF), and the Cascade Asymmetric Linear Filter (CLF), which are two linear variants developed by Dagum and Luati (2009), are evaluated and compared to the 13-term Henderson filter (Henderson, 1916) according to three criteria.

I. The number of unwanted ripples.
II. The size of the revisions made to the most recent estimates when new observations are added to the series.
III. The time lag to detect turning points.

A good method will not only detect turning points fast but it will also minimise the number of times that a time point is incorrectly detected as a turning point. For instance, Zhang M. et al. (2006) computed the percentage of turning points detected and not detected in the trend estimates to quantify differences between seasonally adjusted and trend estimates.

The NLDF and the CLF have already been proven to lead to smaller revisions and produce less false turning points relative to the standard Henderson procedure in previous studies (Dagum, 1996; Chhab, Morry and Dagum, 1999; Dagum and Luati, 2009) but they have never been compared against each other. We are also interested in determining if the same results apply to series from the UK Retail Sales Index.

In section 2, we describe the current practice on the use of trends in ONS. In section 3 we try to address some of the major concerns on the use of trends. Subsection 3.1 gives some background information on the methods under evaluation. Subsection 3.2 describes the data and the statistical measures used. Subsection 3.3 presents and discusses the results of the empirical study. Finally, some concluding remarks and discussion are presented in section 4.

2. Current Practice in ONS

In 1997 an investigation by ONS (Knowles and Kenny, 1997) led to the development of current ONS policy on the calculation, presentation and appropriateness of the use of trends for monthly series.

In summary, the policy recommends two methods of estimating trends on the basis of their good performance with regards to turning points. Both are based on applying Henderson moving averages to seasonally adjusted data. The noise to signal (I/C) ratio, which represents the ratio of the average absolute change in the irregular and in the trend cycle component, is used to determine the length of the Henderson moving average to be used. It is recommended to use the 13-term Henderson moving average with two stages of outlier correction and an ARIMA (0 2 1)(0 0 1) for relatively smooth series (I/C < 1.8) and the 23-term Henderson moving average with no outlier replacement or ARIMA modelling for more irregular series (I/C > 1.8).

Knowles and Kenny (1997) recommended that the use of trends be decided on a case-by-case basis by the ONS business area responsible for the production of the estimates. They offered some advice to decide on whether it is or is not appropriate to present trends, and suggested that these criteria could be modified to define thresholds, beyond which it would be wise to investigate further before publication.

- Speed of detection of turning points: turning points should be detected, on average, in less than 4 months.
- Proportion of false turning points: there should be no more than one false turning point per year, on average (ie less than 9 per cent of the time).
- Smoothness: if the proportion of variation in the time series attributable to irregular movement, as a proportion of variation due to trend, is less than one then the seasonally adjusted series may be considered smooth enough already, with little scope for adding presentational value by showing trends.
- Seasonally adjusted series published: if no seasonally adjusted series is published for a particular output, then it is recommended that trend estimates are not either.
- Data revisions: if source data for a series are subject to revision then trends will be revised when the underlying data are revised and will be less reliable.

According to Knowles and Kenny (1997), even when published, trend estimates should not be quoted as headline figures and all commentary should be written in the past tense. The trend level should normally be presented as a graph on the front page of the press notice (and within the press notice if appropriate) together with the seasonally adjusted series for the last 15 months. The seasonally adjusted series should be shown using a solid line and the trend should be shown as a thicker solid line with a dashed end to reflect the relative uncertainty of the trend at the end of the series. The length of the dashed part of the line is determined by the number of months, on average, before movements in the trend dominate irregular movements in the series. In some cases, to reflect the uncertainty of the Trend 'Trumpet' graphs are also shown in the background notes of the First Release.

The policy led to the introduction of trends in a limited number of releases at ONS, notably for some Labour Market series (Employment, Unemployment, Inactivity Rate, etc.) and also monthly UK Trade. Recently, trend estimates have been used in special supplementary notes accompanying the Retail Sales release.

Most ONS monthly and quarterly series do not show trends. For monthly series, it is common to use 3-month comparisons, such as 3-months on 3-months earlier or with the same 3-month period a year earlier. This is actually just a simple (3-month moving average) trend. One major concern from some users is that the use of trends at any one "snapshot" may look smooth and stable, but over time, as the estimates evolve, revisions to the trend can be considerable. The end of the series is usually where the interest is very strong, and this is the place where the trend is likely to be most vulnerable. Major revision in the trend can lead to a reversal of interpretation, eg false turning points.

In this paper we try to address the issue of false turning points by using alternative methods to the standard Henderson filters used in X-12-ARIMA for smoothing. These methods are explained in detail in the next section.

3. Short-term trend non-parametric predictors for real time analysis

3.1 Background

Dagum (1996) developed a non linear smoother (NLDF) to improve on the classical 13-term Henderson filter, denoted as H13 in the following. The NLDF consists of three steps:

I. Seasonally adjusting the series in the usual way, for example, using X11, X-11-ARIMA or X-12-ARIMA and using the final seasonally adjusted series modified for extremes (Table E2) as the input for step 2.
II. Extending the input series with ARIMA extrapolations.
III. Running the extended series through the Summary Measures option of X-11-ARIMA with strict sigma limits (0.7 & 1.0) for the identification and replacement of extreme values, and taking the final trend-cycle adjusted for level change outliers (Table D12).

The main purpose of the ARIMA extrapolations is to minimise end point revisions to the trend whereas that of extreme values replacement is to reduce the number of unwanted ripples produced by the H13. Note that Table E2 of the X-11 output consists of the modified seasonally adjusted series, that is, those values in the seasonally adjusted series where the irregulars considered to be extreme outliers (given a weight of zero) are replaced by the corresponding trend-cycle estimates (Table D12 of the X-11 output). This is an attempt to ensure the seasonally adjusted series is more typical of the underlying series level, thus preventing the trend estimates from being dragged up or down.

Studies by Dagum, Chhab and Morry (1996) and Chhab, Morry and Dagum (1999) evaluated the performance of the NLDF on a large, representative sample of monthly Canadian socioeconomic indicators of varying degrees of volatility. Results showed the superior performance of the NLDF with respect to both structural and ARIMA standard parametric trend-cycle models. The NLDF method produced significantly less false turning points (unwanted ripples) and smaller revisions than H13 while it retained the timeliness of identifying turning points in the trend estimates. However, given its non-linearity, the NLDF could not preserve the additive constraint by which the trend of an aggregated variable is equal to the algebraic addition of its component trends. To solve this practical problem, Dagum and Luati (2009) propose two linear approximations of the NLDF. There have been concerns that the NLDF causes a shift in the trend level (Chen, 1999), but this has not proved a problem for Statistics Canada who publishes the percentage movements in the trend.

The ALF approximates the NLDF results by combining: (1) the asymmetric weights of an ARIMA (0 1 1) model with $\theta = 0.40$; (2) the weights of a 5-term weighted and a 7-term non-weighted moving average sequentially applied for noise suppression; and (3) the weights of the H13. The ARIMA model and the parameter values are chosen to minimize the size of revisions and the phase shift. The 5-term and 7-term filters, which are applied to the residuals from a first application of the H13 filter, are known to have the good property of suppressing large amounts of power at the frequency band $w = 0.10$, which is the unwanted ripples frequency; and the 7-term filter also for suppressing most of the noise at frequency bands greater than 0.10.

To further reduce the noise, Dagum and Luati (2009) apply the weights of the final cascade symmetric linear filter to this ALF, and generate the CLF. Hence, the weights of the CLF result from combining the previous set of weights (1)-(3) and the weights of the final cascade symmetric linear filter. In other words, the CLF uses twice the standard H13 since the symmetric cascade linear filter is very close to the latter.

Both the ALF and the CLF offer several advantages with respect to the NLDF in addition to preserving the additive constraint. Firstly, their application is direct and hence, does not require knowledge of ARIMA modelling. Secondly, their properties concerning signal passing and noise suppression can always be compared to those of other linear filters. A detailed discussion of their theoretical properties can be found in Dagum and Luati (2009).

We graph their gain functions to illustrate the effect of these linear filters at different frequencies on the amplitude of a cycle for a monthly series. The theoretical properties of these linear methods cannot be compared with those of the NLDF because the latter method is data dependent. The object of applying any filter is to preserve as much of the strength of those cycles of interest as possible, while removing as many cycles outside the range of interest as possible. Therefore, we want the gain function to be close to 1 for frequencies where $w < 0.06$, which means the method preserves the long-term trend component, and near to 0 for w from 0.10-0.50, which means the method removes the irregular component. Figure 1 shows that the gain functions for the last point ALF and CLF filters dramatically suppress the power at $w < 0.10$. H13 passes about 82 per cent of the short cycles (or ripples), ie 9 and 10 months, whereas ALF and CLF pass about 67 per cent of these cycles. Hence, any non parametric method seem to be better than H13 concerning the unwanted ripples criterion.

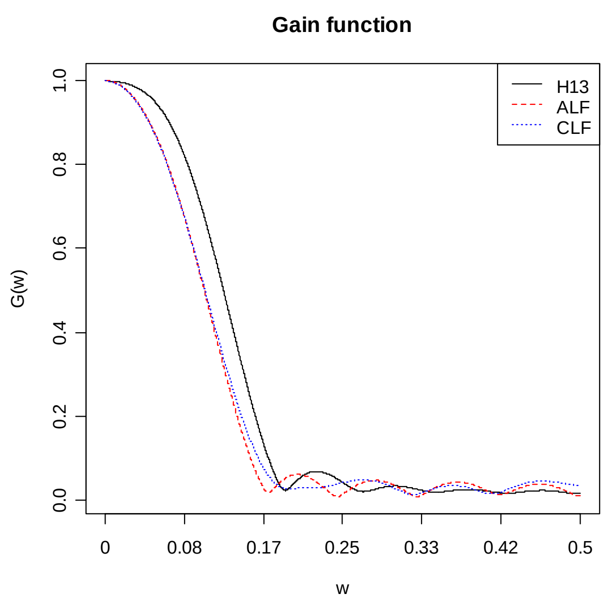

Figure 1: Gain functions of the last point asymmetric filters

The gain function convergence patterns of the asymmetric filters corresponding to H13, ALF and CLF are shown in Figures 2, 3 & 4 respectively. It is evident that the CLF asymmetric filters are very close to one another, and converge faster to the symmetric than H13 or ALF. Due to filtering changes, the asymmetric filters convergence pattern gives an indication of the size of the revisions when new observations are added to the series. Therefore, we expect the CLF to give smaller revisions than H13 and this has been confirmed by Dagum and Luati (2009).

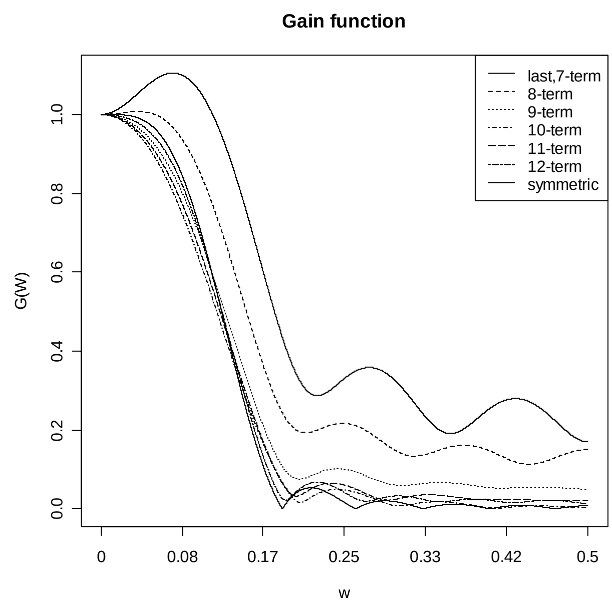

Figure 2: Gain functions convergence pattern of the H13 asymmetric weights to the symmetric

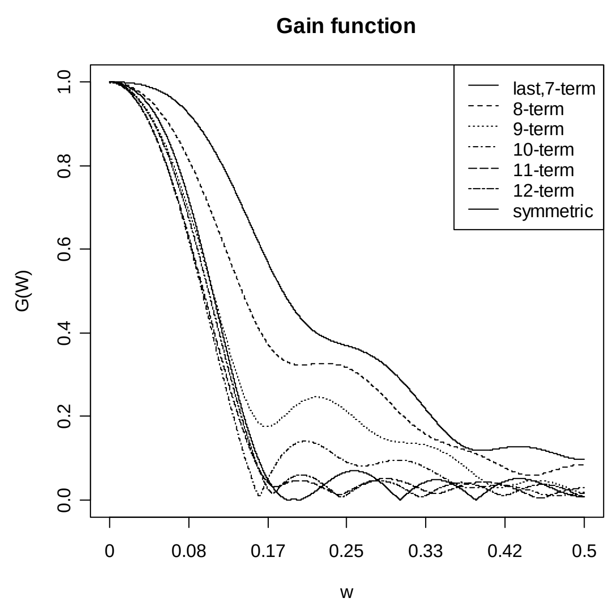

Figure 3: Gain functions convergence pattern of the ALF asymmetric weights to the symmetric

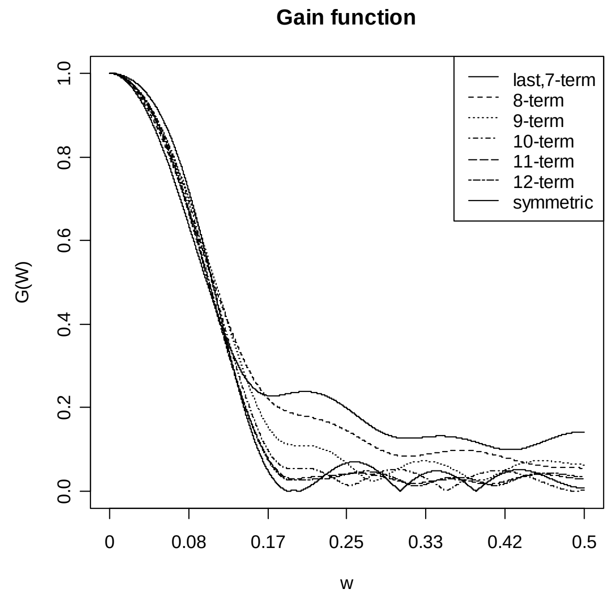

Figure 4: Gain functions convergence pattern of the CLF asymmetric weights to the symmetric

Phase shift properties have not been investigated as part of this study but will form part of future work. However, from Dagum and Luati (2009) it looks like the phase shift for the last point CLF filter is near 2 months (at very low frequencies) relative to H13. Although these differences reduce significantly for the remaining asymmetric filters.

3.2 The data and statistical measures

For the purpose of this study we decided to use the Retail Sales Index because of its ready availability and its importance as a key economic indicator. We used 20 series from the UK Retail Sales Index and 20 series from the Canadian Retail Sales Index. These series had different degrees of volatility based on the noise to signal ratio (I/C) computed in the X-12-ARIMA program (in the range of 1-3.5). The span analysed for all series was from January 1992 to December 2007.

The input to the non-parametric filters was the seasonally adjusted series with extreme values replaced (Table E2 of the X-11 output). Ideally, the ALF and CLF would be applied to the seasonally adjusted data (Table D11 of the X-11 output) since both filters take account of the irregulars in their methodology. However, the D11 and E2 tables for these series were very similar; therefore, for simplicity, we decided to use the same input series as for the NLDF. The trend-cycle estimates given by these non-parametric filters were then compared to the trend-cycle estimate produced by H13 in X-12-ARIMA (Table D12).

To obtain revision measures, the concurrent trend estimates were computed over a four year span: from January 2001 to December 2004 (48 points). By ‘concurrent’ we mean the most recent month's trend estimate by using data up through that month. To produce the first concurrent trend estimate we used the span from January 1992 to January 2001, adding one observation at a time from there on. We then computed the benchmark trend estimate using the symmetric filter on the full span ending in December 2007. The ratio of the absolute percentage difference between the two gives the revision measure for that time point.

The Mean Absolute Percentage Error (MAPE) per series is the average of the 48 revisions, ie:

$$MAPE = \frac{1}{n}\sum_{t=1}^{n}\frac{abs(A_t - B_t)}{B_t}\times 100$$

where

A_t = concurrent trend estimate using the asymmetric filter at time point t
B_t = benchmark trend estimate at time point t

In order to measure the speed of detection of turning points, the dates of the turning points had to be determined first (since the exact dates are unknown for real data series). The key issue of identifying turning points in classical/traditional methods is how to decompose long-term trend and irregular components to extract the cyclical component. In our paper we do not deal with long term filters for detrending seasonally adjusted series but with those that can estimate jointly trend and cycle fluctuations. The main reason for this is that we are concerned with smoothing seasonally adjusted data in the context of rather short series (less than 15 years long) for which their long-term trend is often difficult to identify and estimate accurately. We will therefore apply a generic turning point definition within the context of trend-cycle data.

A turning point is defined as a point in time t when the level of a series, say y_t, is greater/less than the preceding k observations of the series and less/greater than or equal to the subsequent m observations of the series, ie:

$$y_{t-k} \leq \ldots \leq y_{t-1} > y_t \geq y_{t+1} \geq \ldots \geq y_{t+m}$$

defines a downturn at time t and

$$y_{t-k} \geq \ldots \geq y_{t-1} < y_t \leq y_{t+1} \leq \ldots \leq y_{t+m}$$

defines an upturn.

There is no agreement on which values of k and m defines a turning point. Rhoades (1980) defines a turning point for k = 1 and m = 0. Wecker (1979) defines a turning point to be the second of two (or more) successive decreases or increases, ie for k = 2 and m = 2. Zellner, Hong and Min (1991), LeSage (1991) and Pfeffermann and Bleuer (1992) have chosen k = 3 and m = 0. Zhang M. et al. (2006) define a turning point for k = m = 3. For our study, we used the turning point definition for which k = 3 and m = 1 as in Dagum (1996).

A ripple is defined whenever two downturns or upturns occur within a 10 month period (ie small cycles of less than 11 months). The number of unwanted ripples was calculated from the trend-cycle estimate using the full span of the series from 1992 to 2007.

The benchmark turning points were determined by applying Dagum's (1996) definition of turning points to the benchmark trend-cycle estimate. To compute the time lag in identifying the true turning points, we first computed the trend-cycle estimates for all the asymmetric filters of each method (ie the 7-term last asymmetric filter, the 8-term second last asymmetric filter, the third last asymmetric filter and so forth, until the symmetric filters could be used) and then derived the lag by calculating how many months it took for the revised trend series to signal a turning point in the same position as in the final benchmark trend-cycle estimate.

3.3 Empirical comparisons of the methods

Table 1 shows the ratio of the MAPE of the concurrent trend-cycle estimates over a four year period from January 2001 to December 2004. According to these results, any of the non-parametric predictors, regardless of the set of series under study, will give smaller revisions than H13. In particular, both the NLDF and ALF last point asymmetric filters reduce revisions by half.

Table 1: Ratio of MAPE for the last point asymmetric filters relative to the standard H13

Series	NLDF/H13	ALF/H13	CLF/H13
UK Retail Sales	0.52	0.52	0.80
Canada Retail Sales	0.51	0.50	0.89

Table 2 compares the number of unwanted ripples produced by the different trend-cycle symmetric filters. Results suggest that any of the non-parametric predictors will produce, on average, fewer ripples than H13. The ALF and CLF produce the same number of ripples because they both use the same symmetric filter.

Table 2: Number of Unwanted Ripples for the different trend-cycle symmetric filters

Series	H13	NLDF	ALF	CLF
UK Retail Sales	7	4	5	5
Canada Retail Sales	6	2	2	2

Table 3 shows the percentage of turning points detected by each method over the 20 UK and the 20 Canadian time series separately for the period January 2001 to December 2006. Results indicate that the CLF detects the turning points found in the benchmark trend-cycle estimate faster than the rest of the methods. It detects about 46 per cent of the turning points at lag 1 and 66 per cent of the total turning points within two months. In contrast, H13 will only detect about 13 per cent of the turning points at lag 1 and 58 per cent of the total turning points within two months. The NLDF and the ALF also detect a higher percentage of turning points than H13 at lag 1 but they also take longer than four months in a larger number of cases.

Table 3: Lag on Detection of Turning Points (per cent)

Lag	H13	NLDF	ALF	CLF
UK Retail Sales	*nTP**=*119*	*nTP=96*	*nTP=107*	*nTP=107*
1	**14.29**	21.88	20.56	**49.53**
2	43.70	25.00	15.89	15.89
3	30.25	23.96	34.58	11.21
4+	11.76	29.17	28.97	23.36
Canada Retail Sales	*nTP=84*	*nTP=60*	*nTP=72*	*nTP=72*
1	**11.90**	21.67	12.50	**43.06**
2	46.43	28.33	23.61	23.61
3	23.81	23.33	44.44	16.67
4+	17.86	26.67	19.44	16.67

* *nTP* is the number of turning points detected by the symmetric filters over all the series

For illustrative purposes, Figure 5 exhibits the seasonally adjusted series and the trend-cycle estimates of all four methods evaluated for one particular series. Overall, it seems that the CLF reduces the ripples in the trend-cycle data to a greater degree than the other procedures.

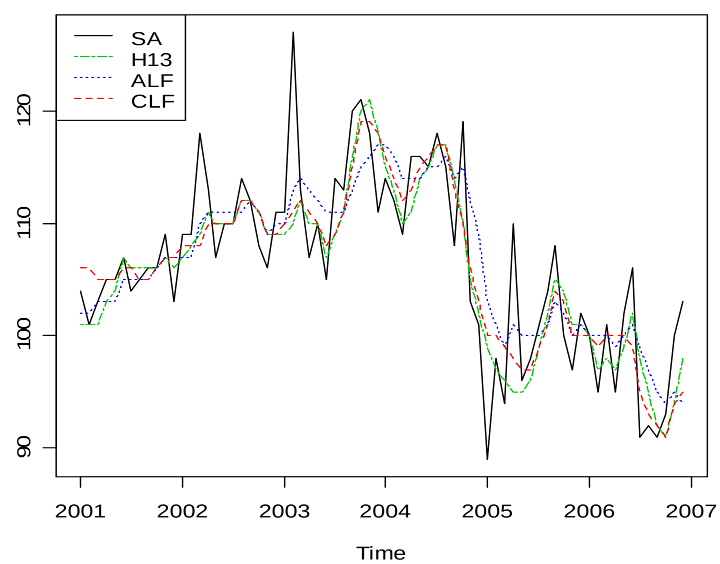

Figure 5: Trend-cycle estimates for the last point asymmetric filters for the span January 2001 to December 2006

4. Conclusions

The analysis confirmed the conclusions from previous studies. The Canadian series were smoother than the UK series but this did not have an impact. Results suggested that either non-parametric predictor will perform better than H13, in particular in terms of revisions and number of unwanted ripples. While the NLDF and ALF reduce by half the size of the revisions to the concurrent trend-cycle estimates, the CLF detects a higher percentage of turning points within two months. The CLF and ALF preserve the additive constraint by which the trend of an aggregated variable is equal to the algebraic addition of its component trends, thus, avoiding the problem of direct versus indirect adjustment.

To summarize, each trend estimation method will produce trends with different properties. Therefore, the most appropriate short-term trend estimation method will depend upon the emphasis given to each of these properties within an NSI.

Acknowledgements

We would like to thank Benoît Quenneville for his help and support. Current work owes much to the unfailing support provided by Pierre Cholette throughout the project.

References

AUSTRALIAN BUREAU OF STATISTICS (2003), ‘A Guide to Interpreting Time Series – Monitoring Trends’, Australian Bureau of Statistics, 1349.0, Canberra, Australia

CHEN L. (1999), ‘Reducing Ripples in ABS trend Series’, A technical report available from the Australian Bureau of Statistics

CHHAB N., MORRY M. and DAGUM E.B. (1999), ‘Further Results on Alternative trend-Cycle Estimators for Current Economic Analysis’, Estadistica, 49-51, 152-157 pp. 231-257

COMPTON S. (2000), ‘Presentation of trend estimates in Official Statistics: UK and international practice’, Proceedings of The Second International Conference on Established Surveys

DAGUM E.B. (1996), ‘A New Method to Reduce Unwanted Ripples and Revisions in trend-cycle Estimates from X-11-ARIMA’, Survey Methodology, Vol. 22, No. 1, pp. 77-83

DAGUM E.B., CHHAB N and MORRY M (1996), ‘Trend-cycle Estimators for Current Economic Analysis’, Proceedings of the Survey Research Methods Section of the American Statistical Association, pp. 140-145

DAGUM E.B. and LUATI A. (2009), ‘A cascade Linear Filter to Reduce Revisions and False Turning Points for Real Time trend-Cycle Estimation’, Econometric Reviews, Vol. 28 (1-3) pp. 4059

FINDLEY D., MONSELL B., OTTO M., BELL W. and PUGH M. (1998), ‘New capabilities and methods of the X-12-ARIMA seasonal adjustment program’, Journal of Business and Economic Statistics 16, pp. 127-177

GRAY, A. and THOMSON, P. (1996), ‘Design of Moving-Average Trend Filters using Fidelity, Smoothness and Minimum Revisions Criteria’, Bureau of the Census, RR96/01

HENDERSON R. (1916), ‘Note on graduation by adjusted average’, Transactions of the American Society of Actuaries, 17, 43-48

KNOWLES J. and KENNY P. (1997), ‘An investigation of trend estimation methods’, A technical report available from Methods and Quality Division, UK Office for National Statistics

LeSAGE, J.P. (1991), ‘Analysis and development of leading indicators using a Bayesian turning-points approach’, Journal of Business and Economic Statistics, 9, pp. 305-316

PFEFFERMANN, D. and BLEUER, S.R. (1992), ‘Probabilistic detection of nonseasonal turning points in economic time series estimated from sample surveys’, Internal report, Methodology Branch, Statistics Canada, Ottawa

RHOADES, D. (1980), ‘Converting timeliness into reliability in economic time series or minimum phase shift filtering of economic time series’, Canadian Statistical Review, 6-13

WECKER, W. (1979), ‘Predicting the turning points of a series’, Journal of Business, 52, pp. 35-50

ZELLNER, A., HONG, C. and MIN, C. (1991), ‘Forecasting turning points in international output growth rates using Bayesian exponentially weighted autoregression’, time-varying parameter, and pooling techniques, Journal of Econometrics, 48, pp. 275-304

ZHANG M., SANDEN N., MENEZES Z. and McLAREN C. (2006), ‘Some Aspects of Turning Point Detection in Seasonally Adjusted and Trend Estimates’, Research Paper, Australian Bureau of Statistics