Comment

A.G. Gray[1] *and P.J. Thomson*[2]

1. Introduction

STL is a simple robust seasonal-trend decomposition procedure built around the loess smoother. It is computationally efficient and is straightforward to use with most of the smoothing parameters selected via simple rules or diagnostic plots. We commend the authors for their simple design and thorough analysis of the filters involved. It is hoped that this study will lead to similar analyses of the filters used in other procedures.

In the discussion that follows we comment on the effects of the loess filters at the ends of the time series, the frequency response functions involved and the iterative analysis. Also included is an analysis of New Zealand building permit data using STL and X-11 together with a concluding section containing general remarks and other observations.

2. End Effects of Loess Filters

Consider a loess filter of nominal length $q = 2r + 1$ applied to an equispaced time series of length $n > q$. (In fact, due to the definition of the neighbourhood weights, the actual length is $q - 2$ in the body and $q - 1$ at the ends of the series.) In the body of the series, loess acts as a symmetric moving average filter with window length q. However, at the ends of the series, its window length remains at q, rather than decreasing to $r + 1$ as in the case of X-11 and other standard procedures. Is this a vice or a virtue?

Conventional wisdom argues that a symmetric smoothing filter used in the body of the series can also, to good effect, be applied at the ends, but with unknown future values replaced by predictions based on past observations. (See Cleveland (1983) and Kenny and Durbin (1982) for example. Dagum (1978) uses this as the basis for the X-11 ARIMA method.) This applies here if the predicted values are generated by loess based on the last q observations and then these observations, augmented by the predictions, are smoothed by a symmetric loess filter of length $2\lambda_q(t) + 1$ centred at a given time point t. To see this, note that if $\hat{g}_t(x)$ is the loess least squares regression curve based on the last q observations and $\lambda_q(t) = s$ then

$$Q = \sum_{-s}^{s} d_j(y_{t+j} - g_t(j))^2$$

$$\geqslant \sum_{-s}^{n-t} d_j(y_{t+j} - \hat{g}_t(j))^2$$

where the d_j are the symmetric loess weights. Thus, if the $y_{t+j} (j > n - t)$ are replaced by $\hat{g}_t(j)$, then Q is minimised when $g_t(j) = \hat{g}_t(j) \quad (-s \leqslant j \leqslant n - t)$ yielding the required loess smooth. Similar considerations apply at the beginning of the series.

[1] On study leave from the New Zealand Department of Statistics.
[2] Institute of Statistics and Operations Research, Victoria University of Wellington, New Zealand.

Acknowledgements: The authors would like to thank the New Zealand Department of Statistics for providing the data used in this commentary. The first author would also like to thank the New Zealand Government Statistician for providing him with leave during part of which this work was undertaken.

The import of the previous paragraph is that the effective window length at the ends of the series may be regarded as $2\lambda_q(t) + 1$. However, over the $r = (q - 1)/2$ observations at each end of the series, $\lambda_q(t)$ increases from r through to $q - 1$ at the extreme ends. Thus the effective window length at the ends is almost twice that in the body of the series. This, in turn, would imply heavier than expected smoothing at the ends of the series than in the body.

To investigate the smoothness properties of STL filters at the ends of the series the following criterion was considered. If the smoothed series is given by

$$m(t) = \sum_{j=1}^{n} w_j(t)y_j(t = 1, \ldots, n)$$

and the $w_j(t)$ are the filter weights, then the smoothness of $m(t)$ was measured by $E\{(\Delta^2 m(t))^2\}$ where Δ^2 is the centred second difference operator. Suppose that y_t follows an additive locally linear model with uncorrelated errors and error variance σ^2. Then, since STL filters pass linear trends,

$$\phi(t) = E\{(\Delta^2 m(t))^2\} = \sigma^2 \sum_{j=1}^{n} \{w_j(t + 1) - 2w_j(t) + w_j(t - 1)\}^2.$$

This criterion should provide a reasonably accurate measure of the smoothness of $m(t)$. Note that $\phi(t)$ is constant, ϕ_0 say, in the body of the series and changes only at the ends. Plots of the relative smoothness $\phi(t)/\phi_0$ over the most recent time points are given in Figure 1, not only for STL filters, but also for the Henderson filters used in X-II. Note, however, that the Henderson filters only pass linear trends approximately at the ends of the series with the approximation worst at the extreme ends. The results reinforce the view that STL significantly smooths the ends of the time series by comparison to the body.

We regard this as a defect, particularly in the case of economic time series where turning points are important to identify. The artificial smoothness at the ends of a series smoothed by STL, particularly over the most recent time points, could lead to unduly optimistic (or pessimistic) assessments of the most recent and future trend values. Practical evidence of this can be seen, to some extent, in the examples given in the paper and also in our example given below. We recommend the adoption of conventional time series windows where the window length decreases from q to $r + 1$. Indeed, in this case the relative smoothness presents a more satisfactory picture as is borne out in Figure 1.

3. Frequency Response Functions and the Iterative Analysis

The operator matrices of the STL filters are near Toeplitz and I–L effectively removes trend. Thus a conventional linear filtering analysis would seem to apply, particularly if end effects are to be ignored. Adopting this approach the transfer functions of the seasonal and trend components after the kth iteration are given by

$$S_k^*(f) = S^*(f)(1 - T_{k-1}^*(f)) \qquad (k = 1, \ldots)$$

$$T_k^*(f) = T^*(f)(1 - S_k^*(f))$$

where $T_0^*(f) = 0$. To ensure that

$$T_k^*(f) = T^*(f)(1 - S^*(f)) + T^*(f)S^*(f)T_{k-1}^*(f)$$

converges, it is essential that $|T^*(f)S^*(f)| \leqslant 1$. However, although the symmetric loess filters are real and bounded in modulus by

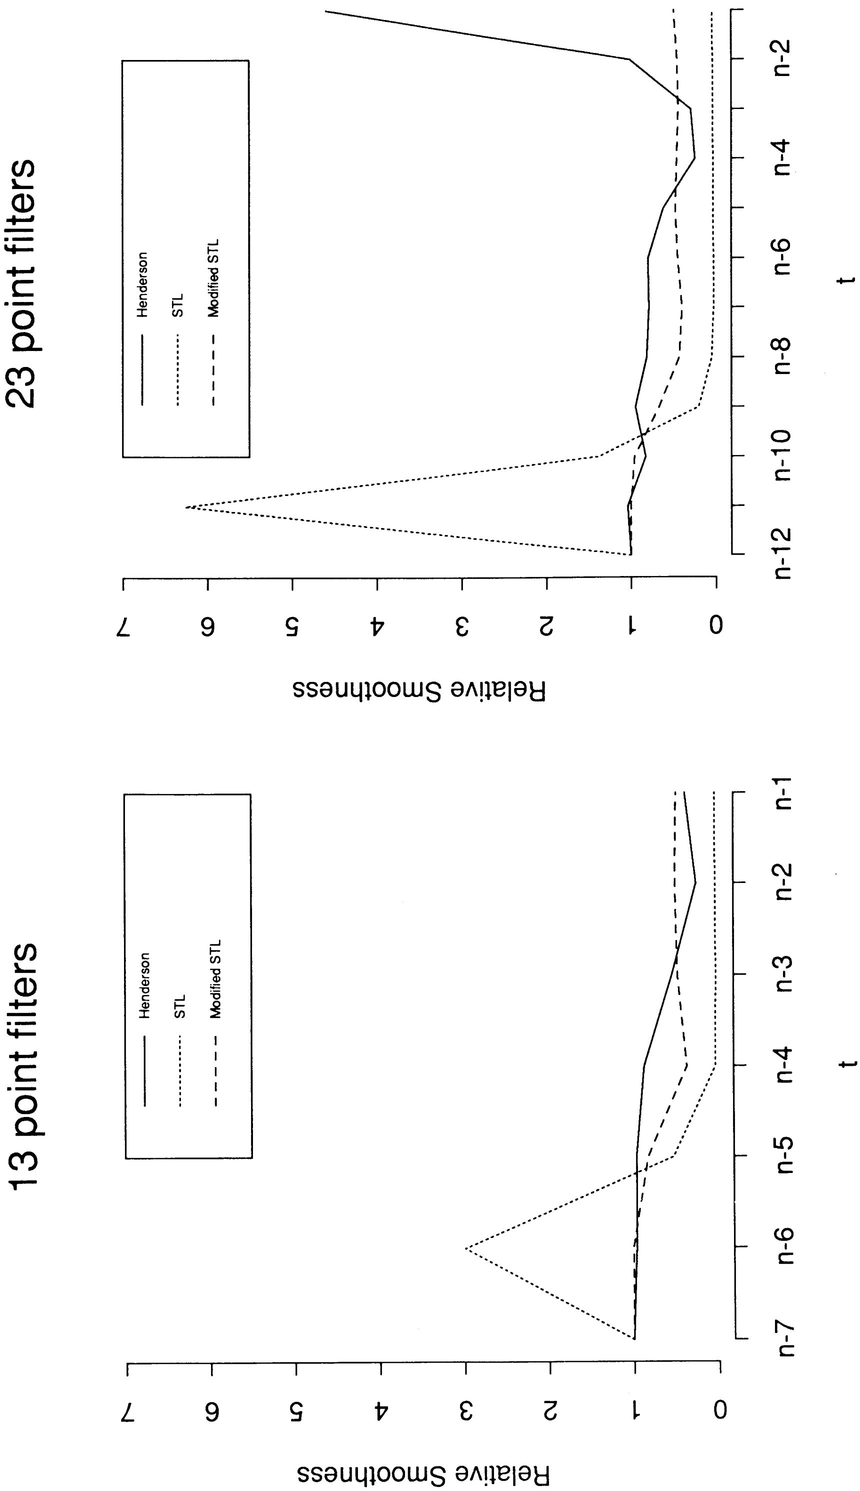

Fig. 1. Smoothness of Filtered Series.

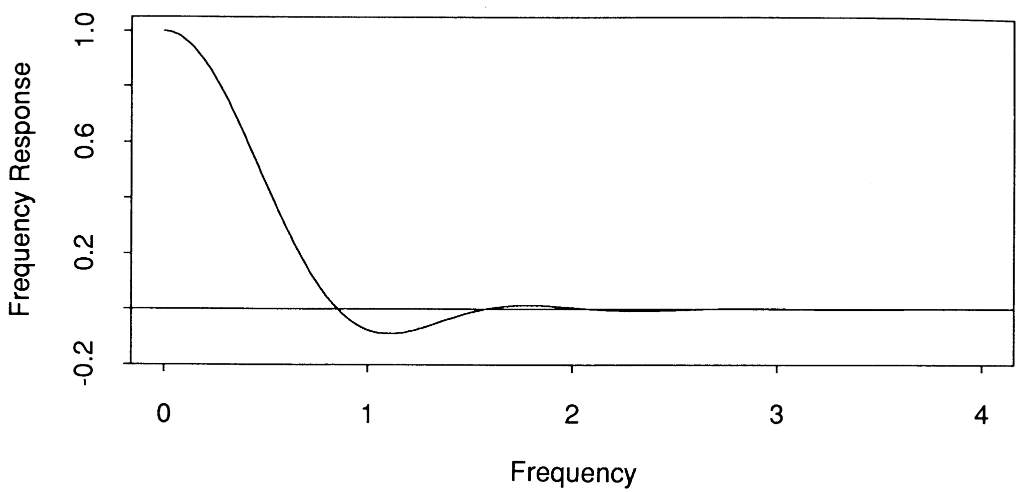

Fig. 2. Limiting Frequency Response Function $H(f)$.

unity, they are not non-negative definite since they can take negative values. Indeed, a straightforward limiting argument shows that $D^*(f)$ defined in §5.2 can be well approximated for r large by $H(rf)$ where

$$H(f) = \frac{\int_0^1 w(u) \cos u\lambda \, du}{\int_0^1 w(u) \, du}.$$

A graph of $H(f)$ is given in Figure 2. Note that the first side lobe is negative. This may prove to be a problem in some circumstances and, as pointed out in §5.4, this is the reason $1 - L^*(f)$ can exceed unity. In the latter case, the choice of $n_{(l)} = [n_{(p)}]_{\text{odd}}$ effectively eliminates the negative side lobe; indeed this choice for $n_{(l)}$ could well be justified in these terms alone. With these caveats $|T^*(f)S^*(f)| \leqslant 1$ and, if $|T^*(f)S^*(f)| < 1$, the transfer functions converge to

$$\tilde{T}^*(f) = T^*(f)(1 - S^*(f))/ (1 - T^*(f)S^*(f))$$

$$\tilde{S}^*(f) = S^*(f)(1 - T^*(f))/ (1 - T^*(f)S^*(f)).$$

Moreover $|T^*(f)S^*(f)|$ can now only equal one when $T^*(f) = S^*(f) = 1$. These observations provide an alternative justification to the analysis given in §5.3.

The negative side lobe problem in $1 - L^*(f)$ could be removed by iterating the loess smoother or replacing the tricube weight function by a non-negative definite weight function. Why was the tricube function chosen? If the tricube function were replaced by another weight function, what smoothness and fidelity criteria should it satisfy? The bisquare function, although not non-negative definite, is equivalent for large q to the optimum filter that passes linear trends and whose second differences have minimum mean square.

4. An Analysis of New Zealand Building Permits Data

The example of the CO_2 data given in the paper is atypical of many economic time series in that it has stable variance, strong additive seasonality and is generally well behaved. By comparison, a typical example of an economic series that the New Zealand Department of Statistics adjusts is the Building Permits for Houses and Flats which is characterized by changing variance, several turning points, large irregulars, and fairly weak seasonality. This monthly data was analysed over the period January 1981 to April 1989 and, since STL does not have a calendar adjustment procedure, was prior adjusted using X-11 to remove trading day effects. A plot of the data is given in Figure 3. Note that the series is too noisy to be able to use Statistics Canada's ARIMA extension of X-11 and too short to apply the sliding spans analysis recently developed by the U.S. Bureau of the Census.

For X-11 we used an additive model and based our choice of seasonal and trend filters on standard X-11 diagnostics. For the seasonal we chose an 11 term (3 × 9) weighted moving average and for the trend we chose a 13 term Henderson weighted moving average. The default values were used for the exclusion of outliers. Note that, since there are at most nine observations for any monthly subseries, nowhere will symmetric seasonal weights be used. The decomposition plot, the cycle-subseries plot, the seasonal-diagnostic plot and the trend-diagnostic plot all indicate that this choice of parameters is appropriate. In this context another useful diagnostic plot is a graph of the original series together with the fitted values plotted by month in much the same way as the seasonal-diagnostic plot.

For STL we followed the guidelines given in the paper and chose a seasonal filter of length 11 and a trend filter of length 21 = [1.5*12/(1 − 1.5/11)]. We also chose to iterate the inner loop once and the outer loop three times. The former choice follows the recommendation of the paper; the latter so that it fits into the X-11 framework. If the convergence criterion given in the paper is adopted then nine iterations are required. However there was little difference between the third and ninth iterations except for June 1987 and July 1988 where the additional

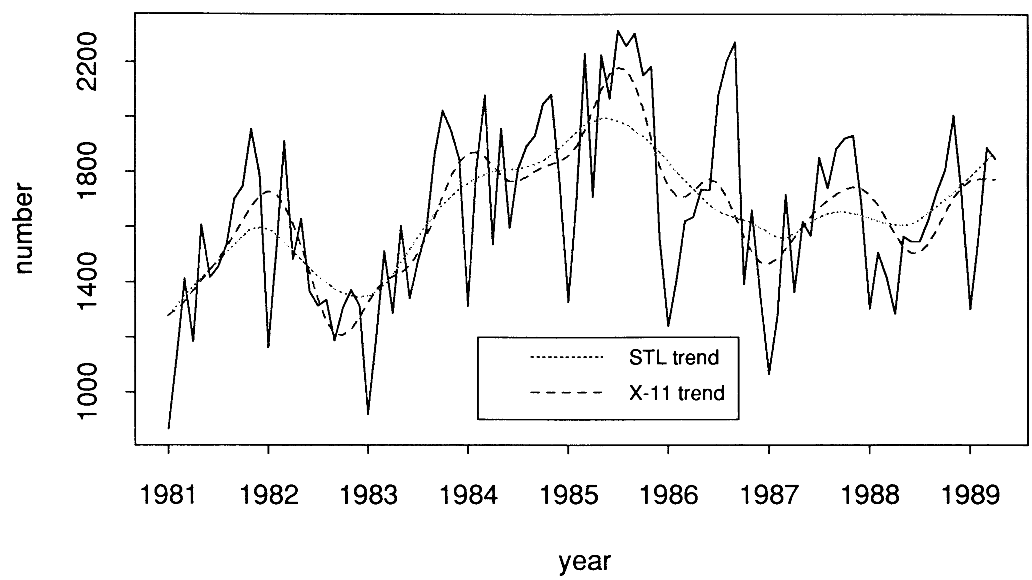

Fig. 3. Building Permits for Houses and Flats.

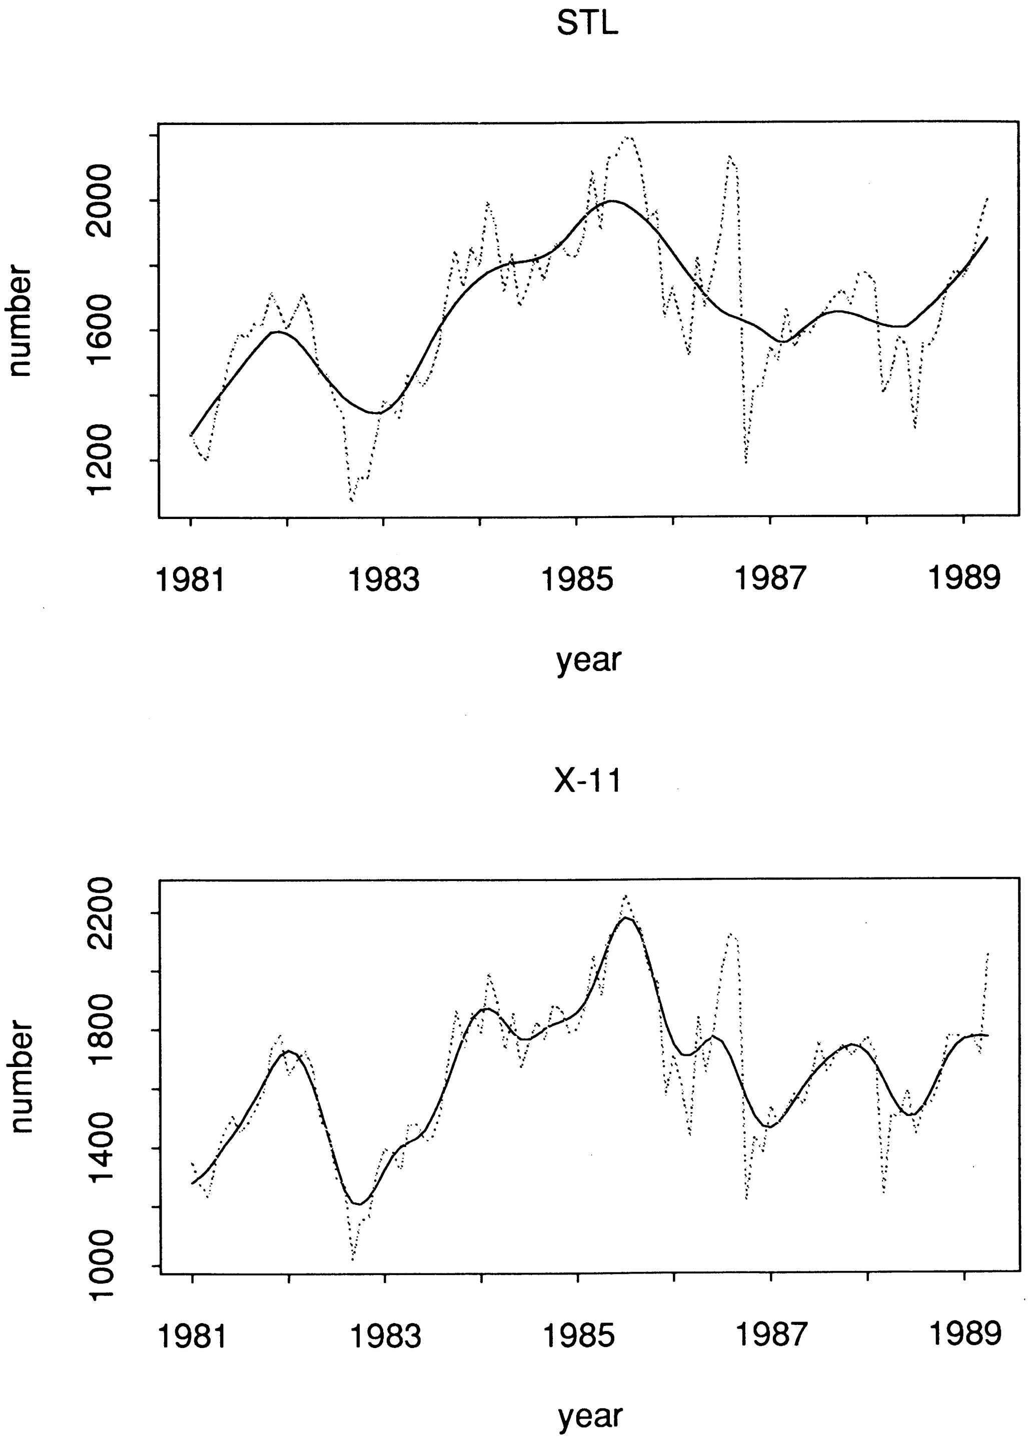

Fig. 4. Building Permits for Houses and Flats; Seasonally Adjusted and Trend.

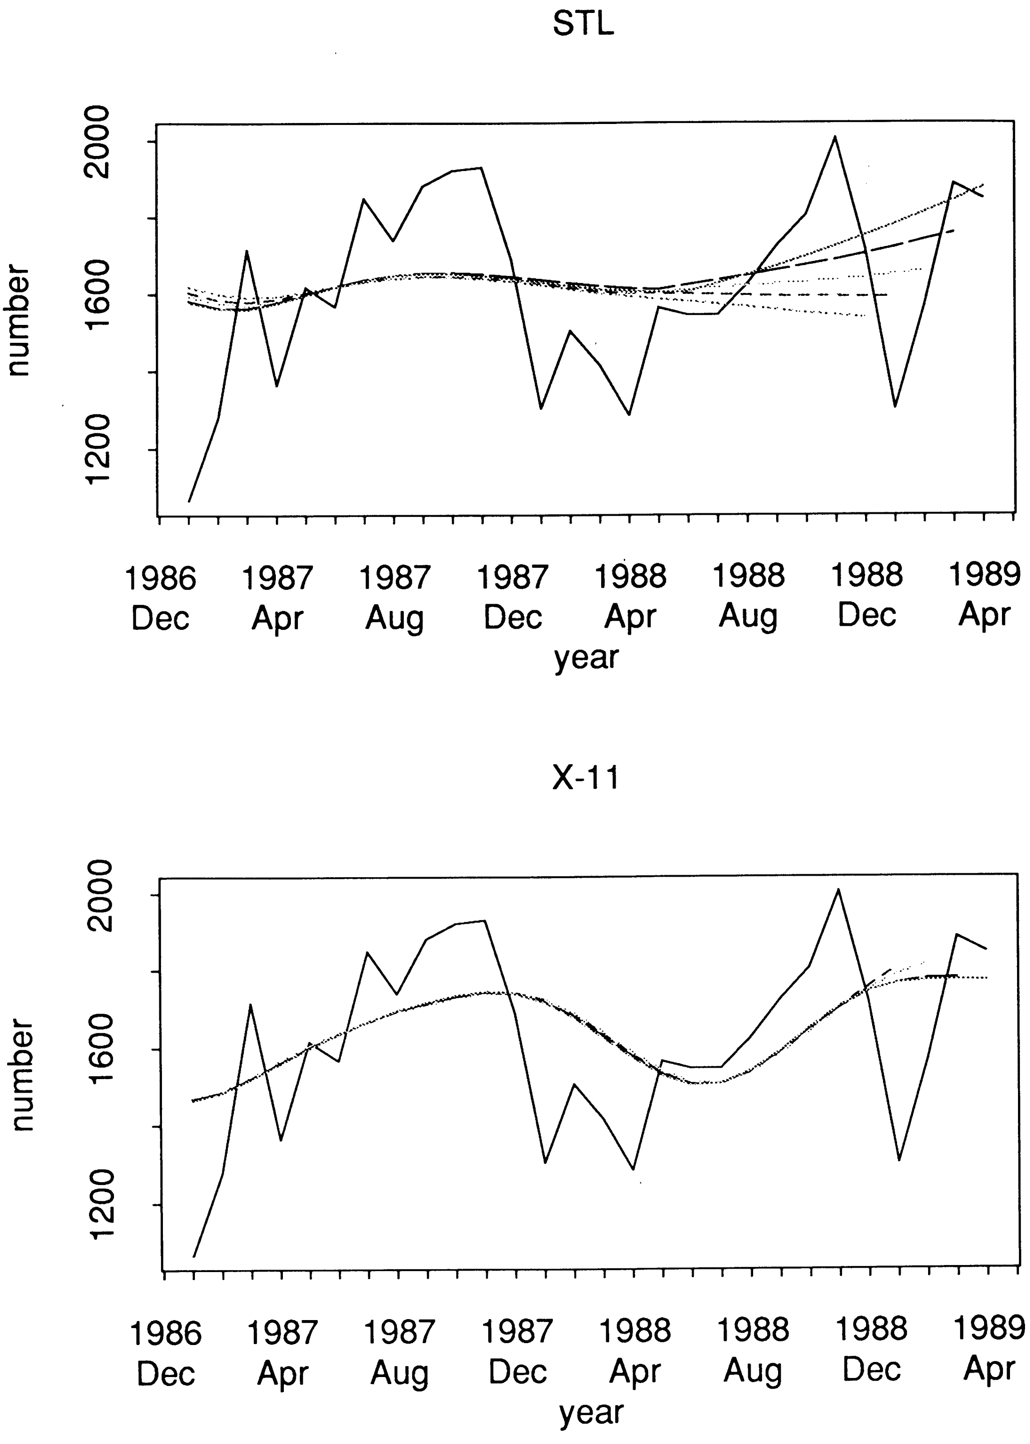

Fig. 5. Building Permits for Houses and Flats; Trend as Data is Added.

iterations strengthened the growth in the seasonal component. Again the various diagnostic plots proved to be satisfactory.

Overall the two procedures gave comparable fits to the data and comparable seasonally adjusted series. However STL fitted the smoother trend and X-11 the smoother seasonal. This was true even when X-11 used the 23 term Henderson moving average to estimate the trend. For this data the balance struck by X-11 seems preferable since its trend better fits the seasonally adjusted data (see Figure 4) and its seasonal is more stable. X-11 also identifies a turning point at the end of the series. The heavier smoothing by STL at the end of the series reduces its sensitivity to turning points and, conversely, makes the near-linear trends at the ends of the series far too sensitive to modest changes induced by the addition of new data. This is illustrated in Figure 5. In the case of STL note the apparent discontinuity in the slope of the trend $(q - 1)/2 + 1 = 11$ points from the end. This can also be inferred from Figure 1. Similar effects were observed with the STL seasonal which typically exhibited more extreme and more volatile near-linear behaviour than X-11, especially at the ends of the series.

5. General Remarks and Conclusions

The simplicity of the design adopted for STL has clearly focussed attention on the key elements involved in the seasonal-trend decomposition of a time series. This is most welcome. However, as a consequence of this clarity, the paper also raises as many questions as it has provided answers, especially with regard to the issues of robustness and end effects.

The theoretical analysis in the paper has concentrated on the effects of STL filters in the body of the series and in the absence of robustness weights. What can be said about the convergence properties of the robustness iterations? How do the robustness weights affect convergence of the inner loop and the properties of the STL filters? Is cycling a possibility?

We agree with the authors that the correct identification and estimation of the seasonal component is the most important step; otherwise trend estimation, whether by eye or by more formal means, will be unreliable. We are not convinced that STL is estimating the seasonal and trend components reliably at the end of the series. With economic time series, where there is noisy data and changing seasonality, the accurate estimation of the components of the decomposition at the ends of the series is of primary importance.

We congratulate the authors for a stimulating and interesting paper which has addressed a number of important practical and theoretical issues in seasonal adjustment and provided a framework for the assessment of others.

6. References

Cleveland, W.S. (1983). Seasonal and Calendar Adjustment. In Handbook of Statistics, eds. D.R. Brillinger and P.R. Krishnaiah, Amsterdam: North Holland, 3, 39–72.

Dagum, E.B. (1978). Modelling, Forecasting, and Seasonally Adjusting Economic Time Series with the X-11 ARIMA Method. The Statistician, 27, 203–216.

Kenny, P.B. and Durbin, J. (1982). Local Trend Estimation and Seasonal Adjustment of Economic and Social Time Series. Journal of the Royal Statistical Society, ser. A, 145, 1–41.