

A REVIEW OF TSW: THE WINDOWS VERSION OF THE TRAMO–SEATS PROGRAM

by D.S.G. POLLOCK

Queen Mary College, University of London,
London E1 4NS, United Kingdom.

Introduction

TRAMO (Time Series Regression with ARIMA Noise, Missing Observations and Outliers) and SEATS (Signal Extraction in ARIMA Time Series) are two closely coupled programs of which the latest versions, that reside within the TSW (TRAMO–SEATS for Windows) program, have become virtually fused.

The TSW program, of which there is both a single-user version and a networked version, is freely available from the website (<http://www.bde.es>) of the Bank of Spain. It is accompanied by a reference manual, which is illustrated with numerous screen shots and which, currently, has 37 pages. The manual contains instructions for the program's installation, instructions for its use, and a glossary and index of the program's input parameters.

The TSW program has been implemented by Gianluca Caporello working under the supervision of Augustin Maravall. The original inspiration for the SEATS program was the work of Burman [2] in the 1980's who built a program for the seasonal adjustment of economic time series while working for the Bank of England. The previous versions of the TRAMO and SEATS programs, which appeared in 1996, were the joint work of Victor Gomez and Augustin Maravall.

It is fortuitous that TSW has reached its present state of maturity at the same time as the Windows version of the rival STAMP program of Koopman *et al.* [7] has been released. Although the two programs share the same objectives, they are underpinned by different philosophies, and they obtain their results in quite different ways. We shall highlight some of these differences in the course of this review.

The essential objective of TSW is to estimate a seasonal ARIMA model and to decompose it into additive components. The estimation is accomplished by TRAMO and the decomposition is by SEATS. This review will concentrate primarily on SEATS, which is where most of the distinctive characteristics of TSW are to be found.

The three regular components into which SEATS resolves a univariate time series model are regarded as the products of ARIMA processes, which are driven by statistically independent white-noise processes.

The *trend-cycle component* is attributed to a model in which the autoregressive (AR) factor comprises a power of the difference operator $\nabla = 1 - L$ together with other factors, if any, that have real-valued positive roots that are close to unity. This accounts for the nonstationary trajectory of the series.

The *seasonal component* is attributed to a model in which the AR factor comprises the seasonal sum $S = 1 + L + \dots + L^{s-1}$, where s is the number of periods in a seasonal cycle, together with other factors containing roots with

THE TRAMO–SEATS PROGRAM

arguments close to the seasonal frequencies. This accounts for the cyclicity of the seasonal process.

The *transitory component* is attributed to a model that comprises complex autoregressive roots with arguments at non-seasonal frequencies.

Finally, the *irregular component* is attributed to a white-noise process. All or part of this component is extracted from the foregoing regular ARIMA components in a manner which ensures that each of them has a zero (i.e. a moving-average root) somewhere on the unit circle.

The inclusion of an irregular component, which results in the so-called canonical decomposition of the univariate ARIMA model, ensures that the regular components of the model are as smooth as possible; which is to say that none of them is roughened by the inclusion of a white-noise element. (The principle of canonical decompositions has been expounded, for example, by Hillmer at Tiao [5] and by Maravall and Pierce [8], and, more recently, by Kaiser and Maravall [6].)

The TRAMO program, which supplies SEATS with an estimated seasonal model, has provisions for estimating regression components and intervention effects, for interpolating missing data values and for handling several types of outliers. It can also handle a range of ugly calendar effects such as trading-day and Easter effects. It succeeds in automatically detecting and correcting data anomalies and it proceeds to an automatic process of model selection and estimation. In the best of circumstances, it works quickly and silently before passing its products to the SEATS program, which will also work without any intervention on the part of the user.

The Program in Use

The user interface of TSW is well thought-out and neatly executed. On starting the program, the user is confronted by a main window, which gives easy access to the essential operations. At the top of the window is a tool bar that contains an array of buttons under which the operations are subsumed.

The *Series* button enables the user to select the data that will be the subject of the analysis. Clicking the button opens the standard Windows dialog box by which the user navigates the file system and selects the data input files.

The *Model* button enables the user to specify a model to be applied to the selected data series. Clicking the button reveals a “tabsheet” of three pages, which allows the user to select the input parameters. The first of these pages presents the user with the choice of whether to accept the automatic modelling procedure or to suspend it in favour of a judicious choice of the input parameters via the two following pages.

The second page enables the user to select the orders of the polynomial lag operators of the seasonal ARIMA model; and it displays the initial values of the parameters, which cannot be altered. It also has a provision for specifying whether or not the data should be mean-corrected and whether logarithms should be taken, albeit that this would be determined automatically by default. Also, the choice can be made between a least-squares estimation and

THE TRAMO–SEATS PROGRAM

a maximum-likelihood estimation of the parameters. The latter is the default option.

The third page of the tabsheet contains the parameters which control the handling of the data anomalies as well as some parameters that govern the more arcane aspects of the automatic processes of model selection. It is here that the regression variables, if there are any, can be added.

The *Run* button sets TRAMO or SEATS or both of them in motion, according to the choice that has been made via the program option panel located on the main window

The *Output* button gives access to some extensive log files relating to the calculations performed by SEATS and TRAMO. The *Output Tables* button reveals the time series that are generated by TRAMO and SEATS in the course of their calculations, whilst the *Out Matrix* button reveals three pages of summary matrices and statistics.

The *Graph* button reveals a window whereby a huge collect of graphs is displayed which represent the primary means of interpreting the combined output of the two programs. The provision is so extensive that the window incorporates its own navigation tree within a full-depth panel. The graphs are filed under four headings, which are ACF (autocorrelation functions), SPECTRA, FILTERS and SERIES.

This profusion of graphs is partly explained by the proliferation of components, which is the consequence of the decompositions. However, the provision goes far beyond what one might reasonably expect; and it represents one of the outstanding features of the program. Nevertheless, there are two omissions, which lead to suggestions for future enhancements.

The first omission is that the graphs have been supplied only in the Windows metafile (.WMF) and bitmap (.BMP) formats. There is no provision for a PostScript output of the sort that generates publication-quality illustrations and which lends itself to manipulation by a skilful user.

The second omission is the absence of empirical periodograms to accompany the model-based spectral density functions of the various components that are calculated using the estimated coefficients. Comparing the periodograms with the estimated spectra could be a useful way of assessing the success of the model in explaining the data and its components.

The Theoretical Underpinnings

The theoretical foundations, which underpin the SEATS program, have been expounded in a number of places by Augustin Maravall and others. The seminal article is that of Burman [2] who expounded the essential algorithms by which the components of the time series are extracted from the data.

We shall recapitulate briefly on these matters; and, in order to simplify the exposition, we shall consider a model from which the transitory component is missing or else fused with the irregular component. In fact, SEATS employs, as its default, the Airline Passenger Model of Box and Jenkins [1] which has no transitory component.

THE TRAMO–SEATS PROGRAM

Consider, therefore, a seasonal autoregressive moving-average model of a sort that may have been estimated by the TRAMO program. This can be written as

$$(1) \quad y(t) = \frac{\theta(L)}{\psi(L)}\varepsilon(t) = \frac{\theta(L)}{\psi_S(L)\psi_T(L)}\varepsilon(t),$$

where $\psi_S(L)$ contains the seasonal autoregressive factors and $\psi_T(L)$ contains the non-seasonal factors.

The denominator polynomial $\psi(L)$ contains both an ordinary differencing operator $(I - L)^d$ and a seasonal differencing operator

$$(1 - L^s)^D = (1 - L)^D(1 + L + \cdots + L^{s-1})^D.$$

The factorisation of seasonal operator implies that the overall degree of differencing within the ARIMA model is $d + D$. The factor $(I - L)^{d+D}$ is assigned to $\psi_T(L)$, whereas $S^D = (1 + L + \cdots + L^{s-1})^D$ belongs to $\psi_S(L)$.

On the assumption that the degree of the moving-average polynomial $\theta(z)$ is at least equal to that of the autoregressive polynomial $\psi(z)$, there is a partial-fraction decomposition of the autocovariance generating function of the model into three components, which correspond to the trend effect, the seasonal effect and an irregular influence. Thus

$$(2) \quad \frac{\theta(z)\theta(z^{-1})}{\psi_S(z)\psi_T(z)\psi_T(z^{-1})\psi_S(z^{-1})} = \frac{Q_T(z)}{\psi_T(z)\psi_T(z^{-1})} + \frac{Q_S(z)}{\psi_S(z)\psi_S(z^{-1})} + R(z).$$

Here, the first two components on the RHS represent proper rational fractions, whereas the final component is an ordinary polynomial. If the degree of the moving-average polynomial is less than that of the autoregressive polynomial, then the irregular component is missing from the decomposition in the first instance.

To obtain the spectral density function of $y(t)$, we set $z = e^{i\omega}$, where $\omega \in [0, \pi]$. (This function is more properly described as a pseudo-spectrum in view of the singularities occasioned by the unit roots in the denominators of the first two components.) The spectral decomposition corresponding to equation (2) can be written as

$$(3) \quad f(\omega) = f(\omega)_T + f(\omega)_S + f(\omega)_R,$$

where $f(\omega) = \theta(e^{i\omega})\theta(e^{-i\omega})/\{\psi(e^{i\omega})\psi(e^{-i\omega})\}$.

Let $\nu_T = \min\{f(\omega)_T\}$ and $\nu_S = \min\{f(\omega)_S\}$. These are the elements of white noise embedded in $f(\omega)_T$ and $f(\omega)_S$. The principle of canonical decomposition is that the white-noise elements should be reassigned to the residual component. On defining

$$(4) \quad \begin{aligned} \gamma_T(z)\gamma_T(z^{-1}) &= Q_T(z) - \nu_T\psi_T(z)\psi_T(z^{-1}), \\ \gamma_S(z)\gamma_S(z^{-1}) &= Q_S(z) - \nu_S\psi_S(z)\psi_S(z^{-1}), \\ \text{and } \rho(z)\rho(z^{-1}) &= R(z) + \nu_T + \nu_S, \end{aligned}$$

THE TRAMO–SEATS PROGRAM

the canonical decomposition of the generating function can be represented by

$$(5) \quad \frac{\theta(z)\theta(z^{-1})}{\psi(z)\psi(z^{-1})} = \frac{\gamma_T(z)\gamma_T(z^{-1})}{\psi_T(z)\psi_T(z^{-1})} + \frac{\gamma_S(z)\gamma_S(z^{-1})}{\psi_S(z)\psi_S(z^{-1})} + \rho(z)\rho(z^{-1}).$$

There are now two improper rational functions on the RHS, which have equal degrees in their numerators and denominators.

According to Wiener–Kolmogorov theory, the optimal signal-extraction filter for the seasonal component is

$$(6) \quad \begin{aligned} \beta_S(L) &= \frac{\gamma_S(L)\gamma_S(L^{-1})}{\psi_S(L)\psi_S(L^{-1})} \times \frac{\psi_S(L)\psi_T(L)\psi_T(L^{-1})\psi_S(L^{-1})}{\theta(L)\theta(L^{-1})} \\ &= \frac{\gamma_S(L)\gamma_S(L^{-1})\psi_T(L)\psi_T(L^{-1})}{\theta(L)\theta(L^{-1})} = \frac{C_S(L)}{\theta(L)\theta(L^{-1})}. \end{aligned}$$

This has the form of the ratio of the autocovariance generating function of the seasonal component to the autocovariance generating function of the process $y(t)$. This formulation presupposes a doubly-infinite data sequence; and it must be translated into a form that can be implemented with finite sequences. There are various alternatives.

The first alternative is to use the Laurent expansion of the rational function on the RHS of (6) as the basis for a moving-average filter (also known as a finite impulse response or FIR filter). Since the Laurent expansion generates a doubly-infinite sequence of coefficients, an approximation based on the central coefficients is called for. This is the method that was employed by Cleveland and Tiao [3] in their implementation of a linearised version of the X-11 seasonal-adjustment filter. The filter is set to work on a data sequence that is supplemented by forecast and backcast values to enable the filter to be run up to the endpoints of the data.

A second alternative, which is followed, for example, by the STAMP program, is to use the Kalman filter and an associated smoothing algorithm to implement a finite-sample version of the filter. Since the data sequence is liable to be nonstationary, a set of nonzero starting values is called for. In the STAMP program, the diffuse Kalman filter algorithm of De Jong [4] is used in handling this aspect of the problem.

A third alternative, which is simpler than using the diffuse Kalman filter, is to use the same filter algorithm in running forwards and backward through the data. This is the finite-sample counterpart of a method that decomposes the filter $\beta_S(L)$ into the forwards filter $\gamma_S(L)\psi_T(L)/\theta(L)$ and the backwards filter $\gamma_S(L^{-1})\psi_T(L^{-1})/\theta(L^{-1})$. As Pollock has shown in [10], the starting-value problem can be handled rather easily in this context.

The fourth alternative, which is the one followed by the SEATS program and which was also Burman’s method, depends upon a partial fraction decomposition of the filter itself. This can be written as

$$(7) \quad \frac{C_S(L)}{\theta(L)\theta(L^{-1})} = \frac{D_S(L)}{\theta(L)} + \frac{D_S(L^{-1})}{\theta(L^{-1})}.$$

THE TRAMO–SEATS PROGRAM

Compared with the preceding approach, this a matter of implementing the filter via components that are joined in parallel rather than in series.

The estimate of the seasonal component obtained by Burman’s method is therefore

$$(8) \quad s(t) = f(t) + b(t) = \frac{D_S(L)}{\theta(L)}y(t) + \frac{D_S(L^{-1})}{\theta(L^{-1})}y(t).$$

This consists of a component $f(t)$, obtained by running forwards through the data, and a component $b(t)$, obtained by running backwards through the data.

In order to compute either of these components, one needs some initial conditions. Consider the recursion running backwards through the data, which is associated with the equation

$$(9) \quad \theta(L^{-1})b(t) = D_S(L^{-1})y(t).$$

This requires some starting values for both $b(t)$ and $y(t)$. The SEATS program obtains these values by stepping outside the sample.

The post-sample values of $y(t)$ are generated in the usual way using a recursion based upon the equation of the ARIMA model, which is

$$(10) \quad \psi(L)y(t) = \theta(L)\varepsilon(t).$$

Here, the requisite post-sample elements of $\varepsilon(t)$ are represented by their zero-valued expectations. The post-sample values of $b(t)$ are calculated by a clever algorithm which was proposed to Burman by Granville Tunnicliffe–Wilson. (Tunnicliffe–Wilson was responsible for writing the programs that accompanied the original edition of the book of Box and Jenkins [1]; and he has played a major role in the development of the computational algorithms of modern time-series analysis.) The Burman–Wilson algorithm is expounded in the appendix to Burman’s paper [2].

To initiate the recursion which generates the sequence $f(t)$, some presample values are found by a method analogous to the one that finds the postsample values.

The Effects of the Distinctive Features of SEATS

The distinctive features of the SEATS program that have been described in the preceding section can have significant effects on the quality of its output.

The first of the distinctive features is the program’s adherence to the principle of canonical decompositions. The effect can be assessed in two ways. First it can be determined how the principle influences the nature of the filters which are used in the program. Then a comparison can be made of the output of SEATS with the output of a program such as STAMP, which does not use canonical filters.

Figure 1 shows, via the solid line, the gain of the trend estimation filter associated with the initial partial-fraction decomposition of the monthly Airline




Figure 1. The gain function of the trend-extraction filter based on a partial-fraction decomposition of the airline passenger model (solid line) together with that of the canonical trend-extraction filter (broken line).

Passenger Model of Box and Jenkins. The indentations in the gain function at the frequencies $\pi j/6; j = 1, \dots, 6$ show how the trend component is deseasonalised. (It can be seen, by analogy with equation (6), that the trend-estimation filter will have the operator $\psi_S(L)$ in its numerator, which means that the poles of the seasonal ARIMA component will become zeros of the filter.)

Disregarding the indentations, one can see how the gain of the filter is reduced only very gradually as the frequency value increases. The trend component extracted by such a non-canonical filter would contain a substantial proportion of the high-frequency components that are present in the original data.

Figure 1 also shows, via the broken line, the gain of the canonical trend-estimation filter that is actually generated by the SEATS program. Here, it can be seen that the imposition of the principle of canonical decomposition has resulted in a radical attenuation of the high-frequency components that enter the trend. However, the estimated trend will continue to comprise components of all frequencies excepting the seasonal frequency of $\pi/6$ and each of the harmonically related frequencies.

The comparison of the gain functions of Figure 1 illustrates one of the major differences between the SEATS program and the rival STAMP program. The STAMP program can generate trends that are noticeably rougher than those generated by SEATS. However, it is remarkable that, in application to the airline passenger data, the two programs generate outputs that are virtually indistinguishable. The reason for this is there is very little power in the airline

THE TRAMO-SEATS PROGRAM




Figure 2. Monthly observations on the number of front-seat passengers killed or seriously injured in road accidents in the UK together with the trend component interpolated by SEATS and (below) the corresponding trend component estimated from the same data by STAMP and displaced downwards.

data in the frequency regions where there are significant differences in the two filters.

Figure 2 shows the trend components of a data series on the road accidents that have been estimated by SEATS and by STAMP respectively. Here the differences are perceptible; and there may be a preference, on aesthetic grounds, for the smoother trend of SEATS. Nevertheless, it is doubtful whether there is any statistical significance in the differences in this case. This is probably the typical case. However, there are cases where the differences are far more marked.

The second distinctive feature of SEATS is its use of the Burman-Wilson algorithm in signal extraction. This methodology is an inheritance from the original program, which Burman wrote for the Bank of England. Nevertheless, is it clear that the retention of the algorithm is a matter of choice. Part of the evidence for this is provided by the presence of the Kalman filter within the TRAMO program, where it is used in evaluating the likelihood function of the ARIMA model. The filter provides an alternative means of extracting the components which could have been implemented easily.

It is interesting to note that, in finding starting values for the signal extraction recursion, it has been decided to extend the sample with forecast and backcast values for $y(t)$ and for the sequences $b(t)$ and $f(t)$, which are present in equation (8). Since there is no need to use forecasts and backcasts of $y(t)$, the inevitable surmise is that stepping outside the sample is a means of tying down

THE TRAMO–SEATS PROGRAM

the ends of the estimated components more securely than would be possible otherwise.

All of the evidence which has come from using the program suggests that the start-up problem, which can adversely affect the estimated components at both of their ends, is handled very successfully by the SEATS program.

An Appraisal of the Program

The TRAMO–SEATS program is addressed to a set of limited and well-defined objectives. Its aim is to implement a model-based methodology of seasonal adjustment and trend estimation in a manner that requires little intervention on the part of the user. In the process of achieving this objective, the program generates a rich variety of output, which is bound to capture the attention of those who are interested in the accompanying processes as much as in the final outcome.

The previous version of the program is already in widespread use amongst banks and governmental statistical agencies in and outside Europe, including Eurostat and the European Central Bank. As Kaiser and Maravall [7] have noted, there is a tendency, fostered perhaps by globalisation and European integration, towards an increasing degree of methodological homogeneity in the analysis of economic statistics. The TRAMO–SEATS program is vying to become a standard tool of analysis that can be used to compare economic activities across countries.

Since the aims and objectives of the program are specific and limited, it cannot reasonably be expected to serve as a general tool of econometric analysis. It is not particularly well-suited, for example, to the needs of exploratory data analysis. It lacks some of the tools that are needed in a program that is used as an habitual tool of data analysis. However, there is much that can be said in favour of small programs tied to specific aims, particularly if they are as carefully constructed, as accessible and as easy to use as this one is.

The final question concerns the scope of the ARIMA-model-based approach to signal extraction. How widely applicable is it to the task of decomposing economic time series into their constituent parts? This is really a question about the nature of econometric time series and of whether their components can be represented adequately by ARIMA models of low orders.

An ARIMA process is generated by subjecting a white-noise process to a linear filtering operation; and, in common with the white-noise process, it comprises cyclical elements at all frequencies within the available range from 0 up to the Nyquist frequency of π , which is at the limit of observability. There are bound to be regions of the frequency spectrum where the ARIMA components of a model-based decomposition overlap substantially. The difficulty in separating such components, one from another, is reflected in the frequency-response profiles of the corresponding signal extraction filters which, typically, show a very gradual transition from pass band to stop band.

Contrary to the presuppositions of the ARIMA-model-based approach, there is evidence that the components of some econometric time series reside in well-defined frequency bands that have little or no overlap. In the

THE TRAMO–SEATS PROGRAM

most favourable circumstances, the components are separated by spectral dead spaces where there are no elements of any significant power. If the tools are sharp enough, then the components can be extracted without loss or confusion. In such cases, we should use other methods in preference to those of the model-based approach.

References

- [1] Box, G.E.P., and G.M. Jenkins, (1976), *Time Series Analysis: Forecasting and Control, Revised Edition*, Holden Day, San Francisco.
- [2] Burman, J.P., (1980), Seasonal Adjustment by Signal Extraction, *Journal of the Royal Statistical Society, Series A*, **143**, 321–337.
- [3] Cleveland, W.P., and G.C. Tiao, (1976), Decomposition of Seasonal Time Series: A Model for the X-11 Program, *Journal of the American Statistical Association*, **71**, 581–587.
- [4] De Jong, P., (1991), The Diffuse Kalman Filter, *The Annals of Statistics*, **19**, 1073–1083.
- [5] Hillmer, S.C., and G.C. Tiao, (1982), An ARIMA-Model-Based Approach to Seasonal Adjustment, *Journal of the American Statistical Association*, **77**, 63–70.
- [6] Kaiser, Regina and A. Maravall, (2001), *Measuring Business Cycles in Economic Time Series*, Springer Lecture Notes in Statistics, Springer-Verlag, New York.
- [7] Koopman S.J., A.C. Harvey, J.A. Doornik, and N. Shephard, (2000), *STAMP 6.0: Structural Time Series Analyser, Modeller and Predictor*, London: Timberlake Consultants Press.
- [8] Maravall, A., and D.A. Pierce, (1987), A Prototypical Seasonal Adjustment Model, *Journal of Time Series Analysis*, **8**, 177–193.
- [9] Pierce, D.A., (1979), Signal Extraction in Nonstationary Time Series, *The Annals of Statistics*, **6**, 1303–1320.
- [10] Pollock, D.S.G., (2000), Trend Estimation and Detrending via Rational Square Wave Filters, *Journal of Econometrics*, **99**, 317–334.