Automatic Identification of Regression-ARIMA Models with Program TSW (TRAMO-SEATS for Windows)

Agustín Maravall
Research Department
Bank of Spain
Alcalá, 48
28015 Madrid

Abstract

The paper presents an overview of the methodology behind program TSW.

TSW is a Windows interface of updated versions of programs TRAMO (Time series Regression with Arima noise, Missing values, and Outliers) and SEATS (Signal Extraction in ARIMA Time Series). The program estimates a general regression-ARIMA model, and computes forecasts and interpolators for possibly no stationary series, with any sequence of missing observations, and in the presence of outliers. Several types of intervention or regression variables can also be included. The program contains an option for automatic model identification, automatic detection and correction of several types of outliers, and for protesting and estimation of Calendar-based effects.

From the model identified for the observed series, filters that provide the MMSE estimators (and forecasts) of the trend, seasonal, calendar, transitory, and noise components in the series are obtained using signal extraction techniques applied to ARIMA models. The program contains a part on diagnosis and on inference, and an analysis of the properties of the estimators and of the estimation and forecasting errors.

Because the programs can efficiently and reliably handle automatic applications to sets of many thousand series, they are already being used intensively in data producing agencies (perhaps the most widely used application is Seasonal Adjustment). They are also used at policy making institutions, business, and research.

Proceedings 23rd European Conference on Modelling and Simulation ©ECMS Javier Otamendi, Andrzej Bargiela, José Luis Montes, Luis Miguel Doncel Pedrera (Editors)
ISBN: 978-0-9553018-8-9 / ISBN: 978-0-9553018-9-6 (CD)

1. INTRODUCTION

Short-term economic data (say, monthly data) are subject to short-term "noise" that, if not properly discounted, can be seriously misleading when interpreting the underlying medium- or long-term evolution of the variables. Among the different types of noise, seasonal variations play a dominant role, and seasonal adjustment of economic data has indeed a long tradition. At present, every month many thousand series are routinely adjusted by data-producing agencies.

Since the appearance of the US Bureau of the Census (USBC) X11 method, adjustment of data has been overwhelmingly done with a few filters designed a priori to have some desirable property, namely, removal of the series variation associated with seasonal frequencies. However, blind application of the X11 filters presents some shortcomings. Lacking a precise definition of seasonality, one seems bound to accept as such simply the output of the filter. Yet insofar as series with no spectral peak at a seasonal frequency may well display power for that frequency (for example, a white-noise series), X11 would estimate a seasonal component for that frequency. This component would of course be spurious.

Ultimately, the seasonal component (or trend component) that a filter estimates is an estimator of an unobserved component and, as such, is contaminated by an estimation error. How large can this error be expected to be? The answer to this question is crucial for a rigorous assessment of the short-term evolution and, naturally, for proper enforcement of policies based on (explicit or implicit) targeting. For example, assume a monetary policy based on annual targeting of a monetary aggregate growth. Further, assume that in December, a growth of 8% is targeted for the next year, and in January the seasonally adjusted series shows a 10% growth (on annualized terms). Could the 2% difference be explained by the size of the measurement error? In other words, is the difference not significantly different from zero? If it is not, the interest rate instrument should be left unmoved; on the contrary, if the difference turns out to be significant, the interest rate should be increased. Clearly, it is important to avoid reacting to noise and not to the proper signal. Unfortunately, computation of the standard deviation of the seasonally adjusted series estimation error does not have a clear answer in the context of a-priori designed filters.

In the early 80's proposals were made to use Auto Regressive Integrated Moving Average (ARIMA) models, together with Signal Extraction techniques of the Wiener-Kolmogorov type to estimate signals (or unobserved components) such as the seasonally adjusted series or the trend. The ARIMA model for each series would be identified from the stochastic structure of the series, and correct specifications could be assessed with standard diagnostics. Next, a stochastic partial fractions decomposition would yield the ARIMA model for each unobserved component, and hence seasonality would be fully defined through its model. The filter, determined from the series structure, would avoid the danger of spuriousness (and, for example, no seasonality would be extracted from a white-noise series). Finally, the model-based structure could be exploited to provide the desired inference, and in particular, the standard deviation of the signal estimation error.

This model-based method was appealing, but it confronted some serious problems. For large scale application it required very heavy doses of time series analyst resources because there were no automatic model identification (AMI) procedures that were reliable enough. Further, it seemed to require heavy computational resources because the identification and estimation algorithms available were not efficient enough. Finally, for many real series, ARIMA models had to be extended in several directions; some important ones were: the capacity to detect and remove possible outliers, possible calendar effects that are non-periodic and other possible special regression-type effects, and the capacity to handle missing observations. To be able to deal with these extensions, ARIMA models have to be replaced by the so-called Regression-ARIMA models. Of course, outlier and calendar effect detection, as well as interpolation, have to be incorporated to the complete AMI procedure.

In the decade of the 90's, Víctor Gómez and Agustín Maravall completed the Regression-ARIMA model-based methodology and produced a pair of connected programs that successfully enforced the methodology. Moreover, they could efficiently and reliably treat in an automatic mode very large sets of series. In the last ten years use of the programs has become widespread throughout the world, and they are intensively used at economic agencies and institutions, research, business, and –most relevantly- in data producing agencies, where seasonal adjustment is the most widely used application. The two programs are named TRAMO ("Time series Regression with ARIMA noise, Missing values and Outliers") and SEATS ("Signal Extraction in ARIMA Time Series"); they are freely available, together with additional tools and documentation, at the Bank of Spain web site (www.bde.es).

The presentation will center on the application of program TSW to a large set of monthly time series. TSW is a Windows interface of updated versions of programs TRAMO and SEATS.

The program estimates a general Regression-ARIMA model, and computes forecasts and interpolators for possibly nonstationary series, with any sequence of missing observations, and in the presence of outliers. The program contains an option for automatic model identification, automatic detection and correction of several types of outliers, and for pretesting and estimation of Calendar-based effects. Several types of intervention or regression variables can also be included.

Next, the program estimates and forecasts the trend-cycle, seasonal, calendar, transitory and noise components in the series, using signal extraction techniques applied to ARIMA models. The program contains a part on diagnosis and on inference, and an analysis of the properties of the estimators and of the estimation and forecasting errors. The last part of the output is oriented towards its use in short-term economic policy and monitoring.

The paper describes the methodology underlying program TSW and, given that the model-based method relies heavily on the adequacy of the complete AMI procedure, special attention will be given to this feature.

At present, production of seasonally adjusted data by data-producing agencies is overwhelmingly dominated by TRAMO-SEATS and X12-ARIMA (the USBC successor of X11). X12ARIMA is a mixed model/filter design-based method that will eventually lead to X13-ARIMA-SEATS. (X12ARIMA has already incorporated and adapted the TRAMO AMI procedure). One is tempted to conclude that, in what concerns seasonal adjustment, a change in paradigm is taking place; the new paradigm being a model-based one.

2. BRIEF DESCRIPTION OF TSW

The first part of the program corresponds to TRAMO; the second part corresponds to SEATS. In the first part, the program estimates and forecasts regression models with errors that follow (most often non-stationary) ARIMA processes. There may be:

- missing observations in the series,
- contamination by outliers,
- contamination by other special (deterministic) effects.

Important cases of the latter are the trading day (TD) effect, caused by the different distribution of weekdays in different months, and Easter effect (EE), which captures the moving dates of Easter.

If B is the lag operator, such that $B\,x(t) = x(t-1)$,

given the observations $y = [y(t_1), y(t_2), \ldots, y(t_m)]$, where $0 < t_1 < \ldots < t_m$,

the model can be expressed as

$$y(t) = \sum_{i=1}^{n_{out}} \omega_i \, \lambda_i(B) \; d_i(t) + \sum_{i=1}^{n_c} \alpha_i \; cal_i(t) + \sum_{i=1}^{n_{reg}} \beta_i \; reg_i(t) + x(t), \qquad (2.1)$$

where

$d_i(t)$: dummy variable that indicates the position of the i-the outlier,

$\lambda_i(B)$: polynomial in B reflecting the outlier dynamic pattern,

cal_i: calendar-type variable,

reg_i: regression or intervention variable,

$x(t)$: ARIMA error,

ω_i: instant i-th outlier effect,

α_i and β_i: coefficients of the calendar and regression-intervention variables, respectively,

n_{out}, n_c and n_{reg}: total number of variables entering each summation term in (2.1).

In compact notation, (2.1) can be rewritten as

$$y(t) = z'(t)\, b + x(t) \quad , \qquad (2.2)$$

where b is the vector with the ω, α and β coefficients, and $z'(t)$ is the matrix with the columns containing the variables in the three summation terms of (2.1).
Let the ARIMA model for $x(t)$ be

$$\phi(B) \;\; \delta(B) \;\; x(t) = \theta(B) \;\; a(t) \quad , \qquad (2.3)$$

where $a(t)$ is a white-noise $(0, V_a)$ innovation.
In (2.3), $\phi(B)$, $\delta(B)$, and $\theta(B)$ are finite polynomials in B. The first one contains the stationary autoregressive (AR) roots, $\delta(B)$ contains the nonstationary AR roots, and $\theta(B)$ is an invertible moving average (MA) polynomial.
Let s denote the number of observations per year. The polynomials assume the multiplicative form

$$\delta(B) = \nabla^{d} \;\; \nabla_{s}^{d_s} \quad ,$$

$$\phi(B) = \left(1 + \phi_1 B + \ldots + \phi_p B^p\right) \;\; \left(1 + \phi_s \;\; B^s\right) ,$$

$$\theta(B) = \left(1 + \theta_1 B + \ldots + \theta_q B^q\right) \;\; \left(1 + \theta_s B^s\right) ,$$

where $\nabla = 1 - B$ and $\nabla_s = 1 - B^s$.

This model will be refered to as the ARIMA $(p, d, q)(p_s, d_s, q_s)_s$ model.
The model consisting of (2.2) and (2.3) will be called a regression (reg)-ARIMA model.
When used automatically, the program

- tests for the log/level transformation,
- tests for the presence of calendar effects,
- detects and corrects for three types of outliers:
  Additive Outliers (AO), Transitory Changes (TC), and Level Shifts (LS); where
  AO: isolated spike;
  TC: spike that dies off gradually;
  LS: a (permanent) step;
- identifies and estimates by maximum likelihood the reg-ARIMA model,
- interpolates missing values,
- computes forecasts of the preadjustment component $z'(t) \;\; b$ and of the ARIMA series $x(t)$ in (2.2).

Note 1 (Estimation)

Writing the model in compact notation as (2.2), by default, b is concentrated out of the likelihood and estimation is iterative:

- Conditional on b $\rightarrow$ ARIMA (MLE),
- Conditional on ARIMA $\rightarrow$ b (GLS).

Note 2: Interpolators and forecasts are obtained as MMSE estimators

3. SUMMARY OF AUTOMATIC MODEL IDENTIFICATION (AMI) PROCEDURE

Pretest for the Log-level Specification

The test consists of direct comparison of the BICs of the default model in levels and in logs (with a proper correction).

Pretest for Trading Day and Easter Effects

The Test is performed with regressions using the default model for the noise and, if the model is subsequently changed, the test is redone. Because underdetection is easier to handle than overdetection, by default, TRAMO has a slight bias towards underdetection.

A Remark on the use of the default (Airline) model

Pretesting and the starting point of AMI depend heavily on the so-called Airline model, given by (for monthly series)

$$\nabla \nabla_{12} x(t) = (1 + \theta_1 B)\ (1 + \theta_{12} B^{12})\ a(t) \qquad (3.1)$$

Three important reasons justify this choice:

- Many studies have shown it is a model appropiate for many real monthly or quarterly macroeconomic series. (Often one finds that an Airline model fits reasonably well between 40% and 60% of the series in a given set.)
- The Airline model approximates well many other models.
  As an example, consider the deterministic “linear trend plus seasonal dummies“ model

  $$x(t) = \mu_0 + \mu_1 t + \sum_{i=1}^{11} \beta_i d_i(t) + a(t) \ ,$$

  where the μ and β parameters are constants, and $d_i(t)$ is a dummy variable with 1 when t corresponds to the it-h month and 0 otherwise. For a standard series with 100 or 200 observations, if the deterministic model yields a good fit, the same will be true for model (3.1) with θ_1 and θ_{12} very close to -1.
- Excellent “benchmark” model.

The model contains 3 parameters (namely, θ_1 , θ_{12} , and V_a) that can be given an intuitive interpretation:

* θ_1 is related to the stability of the trend-cycle component. (Values close to -1 imply stable trends.)

* θ_{12} is related to the stability of the seasonal component. (Values close to -1 imply stable seasonality.)

* V_a measures the variance of the one period-ahead forecast error. (This error will contain, besides the irregular component, the forecasting errors associated with the trend-cycle and seasonal component.)

Automatic Model Identification in the Presence of Outliers

The algorithm iterates between the following two stages

- Automatic outlier detection and correction
- Automatic model identification

The first model used is the default model.

At each step, the series is corrected for the outliers and other regression effects present at the time, and a new AMI is performed. If the model changes, the automatic detection and correction of outliers is performed again from the beginning.

(a) AUTOMATIC OUTLIER DETECTION AND CORRECTION (AODC)

1. Assume that we know the location (t=T), but not the type of outlier.

For T = T, we compute: $\hat{\omega}_{AO}(T),\ \hat{\omega}_{TC}(T),\ \hat{\omega}_{LS}(T)$, the associated t-values

$$\hat{\tau}_{AO}(T),\ \hat{\tau}_{TC}(T),\ \hat{\tau}_{LS}(T)$$

and $\lambda_T = \max\left\{ \left|\hat{\tau}_{AO}(T)\right| ,\ \left|\hat{\tau}_{TC}(T)\right| ,\ \left|\hat{\tau}_{LS}(T)\right| \right\}$.

Use $\lambda_T > C$ to test for significance, where C is an "a priori" set critical value. (The default value depends on the number of observations.)

2. If we don't know the timing of the outlier, we compute λ_t for t = 1, … , N, and use

$$\lambda_T = \max_t \lambda_t = \left|\hat{\tau}_{tp}(T)\right|$$

If $\lambda_T > C$, there is an outlier of type tp (tp can be AO, TC, LS) at T.

We correct for this outlier, and start the process again to see if there is another outlier.

Outliers are removed one by one, until we obtain $\lambda_T < C$.

Now we proceed to joint GLS estimation of the multiple outliers in the full model. A multiple regression is performed to avoid (as much as possible) masking effects.

(b) AUTOMATIC ARIMA MODEL IDENTIFICATION

Suppose the series $\{x_t\}$ follows the model (2.3). TRAMO proceeds in two steps:

First, it identifies $\delta(B)$ (unit roots). Second, it identifies the ARMA model, i.e, $\phi\ (B)$ and $\theta\ (B)$.

Identification of the Nonstationary polynomial $\delta(B)$

To determine the appropriate differencing of the series, we discard unit root (UR) testing for the following reasons.

When regular and seasonal UR may be present, available tests have low power.

For example, in

$$\nabla x(t) = (1 - .8B)\, a(t), \text{ or}$$

$$\nabla_{12} x(t) = (1 - .8B^{12})\, a(t)\ ,$$

the AR unit root would most likely be rejected due to the large MA root.

Besides, in AMI with AODC, one typically tries thousands of models, where the next try depends on previous results. This implies a serious data mining problem: the size of the test is a function of prior rejections and acceptances and the true size of the test is unknown. We follow an alternative approach:

We decide "a priori", instead of a fictitious size, the following value:

How large the modulus of an AR root should be in order to accept it as 1 (unit root).

For AR and MA roots the criterion is different. Roughly, unit AR roots cause no problem; unit MA roots are avoided (in the model for the observed series).

For an AR root, if the modulus is > .95, it is changed to 1 and the root becomes a unit root.

For an MA root, if the modulus is > .99 it is changed to .99, so as to have invertibility. (With .99 no numerical problems appear.)

To identify AR unit roots, we use some useful results (Tiao, Tsay) on superconvengence of unit AR roots. The following example illustrates these results.

Let the true model be the AR (3) model

$$(1 - .5B)\, \nabla^2 x(t) = a(t)$$

Because of superconsistency of UR estimators:

- If we estimate simply an AR(1), the first UR (∇) is likely to be captured.
- If we estimate, again, an AR(1) on the previous residuals, the second UR (∇) is likely to be captured.
- Alternatively, if we start by estimating an AR(2), both UR (∇^2) are likely to be captured simultaneously.
- Further increases in p (the order of the AR), or further AR(1) fits to residuals will not point to a UR.
- The previous results extend in a straightforward manner to seasonal UR.

TRAMO uses these results. First, the model AR (2) $AR_s(1)$ with mean,

$$(1 + \phi_1 B + \phi_2 B^2)(1 + \phi_s B^s)(x(t) - \mu) = (1 + \theta B)(1 + \theta_s B^s) a_t \quad ,$$

is estimated and UR are detected in the way described.

As already mentioned, if the modulus of an MA root is relatively large, the bias in the estimator of the AR parameter can be large, and the U.R. can be missed.

Therefore, after detecting UR with AR fits, TRAMO uses ARMA (1,1) fits to detect UR that might not have been captured because of ignoring possibly large MA roots.

Hence, after the pure AR fit, TRAMO fits multiplicative models of the form ARMA (1,1) $ARMA_s$ (1,1) with mean

$$(1 + \phi B)\left(1 + \phi_s B^s\right)\left(x(t) - \mu\right) = (1 + \theta B)\left(1 + \theta_s B^s\right)a(t) \quad .$$

Unit roots are identified 1 by 1.

The residuals of the last estimated model are used for a pre-test to specify a mean or not.

Identification of the stationary ARMA model: $\phi(B)$ *and* $\theta(B)$

The program selects the orders (p,q), where $p = dg\{\phi(B)\}$ and $q = dg\{\theta(B)\}$, corresponding to the lowest BIC (p,q), where

$$BIC\ (p,q) = \ln\left(V_a\ (p.q)\right) + (p + q)\frac{\ln(N - d)}{N - d},$$

and V_a is the residual variance associated with that particular model and d the number of observations lost by differencing.

It searches for models of the form

$$\phi_{p_r}(B)\,\phi_{p_s}\left(B^s\right)\tilde{x}_t = \theta_{q_r}(B)\,\theta_{q_s}\left(B^s\right)a_t \quad , \qquad (3.2)$$

where $\tilde{x}(t) = \nabla^d \nabla^{ds} x(t)$ is the differenced series, over the range

$0 \le p_r, q_r \le 3, \qquad 0 \le p_s, q_s \le 2$ (1 if used with SEATS)

This is done sequentially (for fixed regular polynomials, the seasonal ones are obtained, and viceversa).

The search favours parsimony and "balanced" models (similar AR and MA orders). Parsimony improves estimation precision and out-of-sample forecasting performance. Balanced models tend to be stable over time.

Remark:

When analyzing series with care, TRAMO may suggest a few models (perhaps 2 or 3) all of which could be reasonably acceptable.

When used with SEATS, looking for, among these models, the one that provides a more satisfactory decomposition (e.g.,a better seasonal adjustment) may provide additional tools for the choice.

4. DECOMPOSITION OF THE SERIES AND SEASONAL ADJUSTMENT

The first part of the program (associated with TRAMO) estimates the possible outliers, calendar and regression effects, which are treated as deterministic, and hence decomposes the observed series y(t) into a deterministic and a stochastic component (the two terms $z'(t)b$ and x(t) of equation (2.2), respectively). The deterministic component $z'(t)\hat{b}$ is called the “preadjustment” component, and once it is removed from y(t), an estimate of the stochastic part is obtained. This stochastic part is assumed the output of a linear stochastic process (as in equation (2.3),) and is also refered to as the “linearized series.”

In the second part of the program (associated with SEATS) the ARIMA-model-based (AMB) methodology is used to estimate stochastic unobserved components in x (t). The unobserved components are the

-trend cycle, p(t),

-seasonal, s(t),

-transitory, c(t),

-irregular, u(t),

and, for an additive decomposition,

$$x(t) = p(t) + c(t) + u(t) + s(t) = n(t) + s(t) \quad , \qquad (4.1)$$

where n(t) denotes the SA series. (A multiplicative decomposition can be transformed into an additive one by taking logs.)

Broadly, the trend-cycle captures the spectral peak around the zero frequency, the seasonal component captures the peaks around the seasonal frequencies, the irregular component picks up white-noise variation, and the transitory component captures transitory variation that differs from white noise.

From the ARIMA model for the series, the models for the components are derived. This is done by partitioning the spectrum (or pseudo-spectrum) of x(t), as described in Maravall (1995) or Gómez and Maravall (2001); the procedure follows the approach of Burman (1980).

The models for the components also have ARIMA-type expressions. Typically, for the trend-cycle and seasonal component,

$$\nabla^D \; p(t) = w_p(t) , \qquad D = d + d_s \tag{4.2}$$

$$S \; s(t) = w_s(t), \qquad S = 1 + B + \ldots + B^{s-1} , \tag{4.3}$$

where $w_p(t)$ and $w_s(t)$ are stationary ARMA processes. These models are always balanced.

The transitory component is a stationary ARMA process and the irregular component is white noise.

The processes $w_p(t)$, $w_s(t)$, c(t), and u(t) are assumed to be mutually uncorrelated.

Aggregation of the models for p, s, c, and u yields the ARIMA model (4.3) that had been identified for the series x(t).

The model for the SA-series is obtained through aggregation of the models for p(t), c(t), and u(t). Its basic structure is also of the type (4.2), with p replaced by n.

It is well-known that, in (4.1), there are a variety of ways in which the additive white noise can be assigned to the components. Identification is achieved in SEATS by imposing the "canonical condition", whereby all additive white noise is assigned to the irregular component. In this way, the variance of the later is maximized, and the rest of the components are as stable as possible, given the stochastic features of the series.

As an example, for the Airline model (3.1), the model for the trend-cycle is an IMA (4.2), say

$$\nabla^2 \; p(t) = (1 + \theta_{p,1} \; B + \theta_{p,2} \; B^2 \;) \; a_p(t), \tag{4.4}$$

where the MA polynomial can be factorized as

$$\theta_p(B) = (1 + \alpha_p B)(1 + B) \; ,$$

with the root B=-1 inducing a spectral zero for the frequency $\omega = \pi$, associated with the canonical condition. The parameter α_p often takes values that are close to -1.

The model for the SA series is also of the type (4.4) - with "p" replaced by "n" - and the factorization of the MA polynomial yields two real roots, one close to 1, the other, a relatively small number. Thus, the SA series is not far from the "random walk+drift" model.

The model for the seasonal component is as in (4.3), with $w_s(t)$ following an MA (s-1) model. (This MA will contain a unit root for a non-seasonal and non-trend frequency.)

Finally, the irregular component is white noise and there is no transitory component.

Remark: On (relatively rare) occasions, the ARIMA model identified by TRAMO does not accept a decomposition such that all components have, for all frequencies, a non-negative spectrum. It is then said that the model presents a "non-admissible" decomposition. When this happens, SEATS automatically modifies the model, searching for a decomposable model that is not far from the TRAMO one. The search always converges.

The component estimator and forecast are obtained, by means of a Wiener-Kolmogorov type filter, as the MMSE estimators (under the normality assumption, equal to the conditional expectation) of the signal given the observed series.
Although its implementation does not require knowledge of the parametric models for the components, the model-based structure offers a representation that is useful for analytical interpretation.
The filter can be given a simple analytical representation. In more compact notation,
let $x(t)$ follow the model

$$\phi(B) \quad x(t) = \theta(B) \ a(t) , \qquad a(t) \sim wn(0, V_a) , \qquad (4.5)$$

where now $\phi(B)$ also includes the unit roots.
Consider the decomposition of x(t) into "signal plus non-signal" as in

$$x(t) = s(t) + n(t) .$$

Let the model for the signal s(t) be

$$\phi_s(B) \quad s(t) = \theta_s(B) \ a_s(t) , \qquad a_s(t) \sim wn(0, V_s) ,$$

where $\phi_s(B)$ will contain the roots of $\phi(B)$ associated with the component s.
Analogously, let the non-signal, n(t), follow the model

$$\phi_n(B) \ n(t) = \theta_n(B) \ a_n(t), \qquad a_n(t) \sim wn(0, V_n),$$

where $\phi(B) = \phi_s(B)\phi_n(B)$. Then, if $F = B - 1$ denotes the forward operator (such that $F\,x(t) = x(t+1)$,) for a doubly infinite series, the WK filter to estimate the signal is given by

$$\nu_s(B, F) = \frac{V_s}{V_a} \quad \frac{\theta_s(B) \quad \phi_n(B)}{\theta(B)} \quad \frac{\theta_s(F) \quad \phi_n(F)}{\theta(F)} , \qquad (4.6)$$

or, equivalently, by the ACF of the stationary ARMA model

$$\theta(B) \quad z(t) = \left[\theta_s(B) \quad \phi_n(B) \right] b(t) , \qquad b(t) \sim wn \ (0, \ V_s \ / \ V_a) \ .$$

The filter given by (4.6) is two sided, centered, symmetric and convergent.
The estimator of the signal is obtained through

$$\hat{s}(t) = \nu_s \ (B, F) \quad x(t) \ . \tag{4.7}$$

In practice, one deals with a finite series: $[x(1), x(2), \dots , x(T)]$.
Given that the WK filter converges, for long-enough series the estimator of the signal for the mid-years of the sample can be considered to be equal to the historical estimator (that is, the one that would be obtained with the doubly infinite series).
The model based structure allows us to derive formulas that can be exploited for diagnostics and inference.
First, from (4.5) - (4.6) it is obtained that the historical estimator is generated by the model

$$\phi_s(B) \ \hat{s} \ (t) = k_s \ \theta_s \ (B) \ \frac{\theta_s \ (F) \quad \phi_n \ (F)}{\theta \ (F)} \ a_t \quad , \tag{4.8}$$

where $k_s = V_s \ / \ V_a$. Therefore, the theorical variance, autocorrelation generating function (ACF) and spectrum of the stationary transformation of $\hat{s}(t)$ are straightforward to obtain. Using Bartlett´s approximation, it is possible to compute the standard errors (SE) of the empirical estimators of the variance and autocorrelations of the stationary transformation of $\hat{s}(t)$ -trimming some years at both ends of the series - and derive an approximate test for under/over estimation of s(t), and for possible misspecification of the model. Further, it can be seen that the theorical covariance of $\hat{s}(t)$ and $\hat{n}(t)$ (the historical estimator of the non-signal) is equal to the variance of the ARMA model

$$\theta(B) \ z_t \ = \theta_s(B) \ \theta_n(B) \ b(t) \ , \tag{4.9}$$

with Var $b(t) = V_s \ V_n \ / \ V_a$. Because $\theta(B)$ is invertible, (4.9) is a stationary process with finite variance. It follows that the empirical contemporaneous crosscorrelation between $\hat{s}(t)$ and $\hat{n}(t)$ when (4.5) is non-stationary should be relatively small and positive.
The final or “historical” estimator of s(t), to be denoted $\hat{s}(t)$, obtained with a doubly infinite filter, contains an error, equal to $h(t) = s(t) - \hat{s}(t)$, to be denoted “historical estimation error” (HEE). Its variance and ACF turn out to be those of the stationary ARMA process given by (4.9).

Even for non-stationary series with theorically infinite variance, the HEE will have finite variance and therefore the theorical component s(t) and its estimator $\hat{s}(t)$ will be co-integrated.

Notice that, given that in general $\hat{n}(t)$ will be non-stationary, $\phi_n(B)$ -and hence $\phi_n(F)$- will contain unit roots. According to (4.8), $\hat{s}(t)$ will be a non-invertible series. For example, for series with seasonality, the historical estimator of the SA series and of the trend-cycle will typically contain seasonal U.R. in their MA polynomials.

More generally, given the series $[x(1),\dots,x(T)]$, the MMSE estimators and forecasts of the components are obtained applying the two-sided WK filter to the series extended at both ends with forecasts and backcasts. These are computed with (4.5), the ARIMA model for the observed series.

Therefore, for a finite realization, the estimator of s(t) when observations end at period T can be expressed as

$$\hat{s}(t \mid T) = \nu_s(B,F)\ \hat{x}(t \mid T) \tag{4.10}$$

where $\hat{x}(t \mid T)$ denotes the extended series. By extending the series with an appropriate number of forecasts, (4.10) provides, besides in-sample estimators, MMSE forecasts of the signal, i.e., $\hat{s}(t \mid T)$ for $t > T$.

Applying the Burman-Wilson algorithm, it is possible, however, to obtain the full effect of the doubly infinite filter with just a small number of forecasts and backcasts. (For example, 26 are needed to estimate the SA series in a series that follows a monthly Airline model.)

The model-based framework is exploited to provide (SE) of the estimators and forecasts (as well as of the associated rates of growth).

Being obtained by using forecasts, the component estimators at the end points of the series are preliminary, and will suffer revisions as future data become available, until it can be assumed that the historical estimator has been reached. These revisions typically last between 2 and 5 years.

The model-based framework is also exploited to analyze revisions (size and speed of convergence) and to provide further elements of interest to short-term monitoring.

We will maintain the semi-infinite realization assumption for the series. (In McElroy, 2006, it is seen that, for the vast majority of cases, the finite realization assumption yields considerably close results.)

Assume the complete WK filter to estimate the signal for period t can be safely truncated at observation t+K.

If $\hat{s}(t \mid t+k)$ denotes the preliminary estimator, the revision it will suffer is given by

$r_k(t) = \hat{s}(t) - \hat{s}(t \mid t+k)$.

Notice that $r_k(t)$ will decrease as k increases, until it can be assumed equal to 0 for large enough K.

From (2. 8), $\hat{s}(t)$ can be expressed as

$$\hat{s}(t) = \frac{\alpha(B)}{\phi_s(B)} a(t+k) + \frac{\beta(B)}{\theta(B)} a(t+k+1) , \qquad (4.11)$$

where $\beta(B)/\theta(B)$ converges to zero. The polynomials $\alpha(B)$and $\beta(B)$can be easily obtained as in Bell (2007) or Maravall (1994). Taking expectations at time (t+k), the second summation term vanishes, so that the preliminary estimator follows the model

$$\phi_s(B)\ \hat{s}(t \mid t+k) = \alpha(B)\ a(t+k) \qquad (4.12)$$

where B and F will always operate on t. From (4.11) and (4.12), it follows that the revision in $\hat{s}$ (tlt+k) follows in turn the stationary ARMA model

$$\theta(F)\ r(t \mid t+k) = \beta(F)\ a(t+k+1) . \qquad (4.13)$$

The estimation error contained in the estimator (or forecast) $\hat{s}$(tlt+k) can be expressed as

$$e(t \mid t+k) = s(t) - \hat{s}(t \mid t+k) = [s(t) - \hat{s}(t)] + [\hat{s}(t) - \hat{s}(t \mid t+k)] . \qquad (4.14)$$

The first bracket in the r.h.s. of (4.14) is the historical estimation error, h(t), that follows the ARMA model (4.9). The second bracket is the revision error, $r(t \mid t+k)$, that follows model (4.13). The two errors are uncorrelated so that, from

$$e(t \mid t+k) = h(t) + r(t \mid t+k) .$$

The model for the total estimation error is easily obtained. Its variance and ACF are those of the model

$$\theta(B)\ e(t \mid t+k) = \delta(B)\ d(t) , \qquad (4.15)$$

where the MA part in the sum of two (finite) MA´s, i.e.

$$\delta(B)\ d(t) = \theta_s(B)\ \theta_n(B)\ b(t) + \beta(B)\ a(t+k+1).$$

From models (4.9), (4.13), and (4.15), inferences can be drawn concerning errors in the different types of estimators. From the variances, confidence intervals can be computed, and knowledge of the ACF permits us to compute errors in the growth (or linearized rates of growth when logs have been taken) over any desired period.

Three remarks are worth mentioning:

(a) Preliminary estimator, historical estimator and component –i.e., $\hat{s}(t \mid t+k)$, $\hat{s}(t)$, and $s(t)$ - share the same AR polynomial and unit roots.

(b) The total estimation error, the error in the historical estimator, and the revision error share the same stationary AR polynomial, equal to the MA polynomial of the model for the series. Therefore the three errors will be stationary so that preliminary, historical estimator and component are all pairwise cointegrated.

(c) Given that the "true" component $s(t)$ will never be observed, the error in the historical estimator is more of academic rather than practical interest. In practice, interest centers on revisions. From (4.13), the revision standard deviation, will be an indicator of how far we can expect to be from the optimal estimator that will be eventually attained, and the speed of convergence of $\theta(B)^{-1}$ will dictate the speed of convergence of the preliminary estimator to the historical one. (Fortunately, most often, very slow convergence is associated with very small revisions.) This information allows us to answer questions of applied concern, such as for example:

- Is it worth it to move from a once-a-year adjustment to a concurrent one?
- For how long should the series be revised?

5. MIXING THE TRAMO AND SEATS RESULTS

1. <u>*Final Estimators*</u>

The stochastic components estimated by SEATS provide the decomposition of the linearized series that has been preadjusted by TRAMO. The (mostly deterministic) effects estimated by TRAMO that form the preadjustment component can be combined with the SEATS estimators in order to obtain the final estimators. By default,

- Level Shift outliers will be assigned to the final trend.
- Additive and Transitory Change outliers will be assigned to the final irregular component.
- Calendar effects (Trading Day, Easter, Leap Year, moving holidays,…) will go to the final seasonal component after proper centering. The mean resulting from the centering will go the trend. (They also may form a separate component.)
- Regressions and intervention variable effects can form a separate component or be assigned to the one thought appropriate. (If assigned to the seasonal component, the effects will always be centered and the resulting means added for the trend.)

2. <u>*Change of Models*</u>

On occasion, SEATS may change the ARIMA model selected by TRAMO, and use the modified one to decompose the series. The two most important reasons for that to happen are the following.

- The ARIMA model of TRAMO does not accept an admissible decomposition. An admissible decomposition is one in which all the component's spectra are non-negative for all frequencies.

- The model specification is changed in a way that is likely to improve stability of the seasonal and (perhaps) trend-cycle components. This mostly affects highly stationary seasonal structures, which tend to generate highly unstable seasonal components.

Additionally, the model can be changed in order to avoid a nonsense decomposition. An example can be a seasonal ARMA (1.1) structure such as

$$(1 + \phi_{12} B^{12}) \ x(t) = (1 + \theta_{12} B^{12}) \ a(t) \ ,$$

with $-1 < \theta_{12} < \phi_{12} < 0$. For this structure the spectral peaks appear at non-seasonal frequencies. (Nonsense models are very rarely encountered.)

6. REFERENCES

BELL, W.R. (1984), "Signal Extraction for Nonstationary Time Series", *Annals of Statistics* 12,646-664.

BURMAN, J.P. (1980), "Seasonal Adjustment by Signal Extraction", *Journal of the Royal Statistical Society* A, 143, 321-337.

CAPORELLO, G., and MARAVALL, A. (2004), "Program TSW: Revised Reference Manual", Working Paper 0408, Servicio de Estudios, Banco de España.

CAPORELLO, G. and MARAVALL, A. (2003), "A Tool for Quality Control of Time Series Data. Program TERROR", Occasional Paper 0301, Servicio de Estudios, Banco de España.

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

FINDLEY, D. F. and MARTIN, D. E. K. (2006), "Frequency Domain Analyses of SEATS and X-11/12-ARIMA Seasonal Adjustment Filters for Short and Moderate-Length Time Series", *Journal of Official Statistics*, Vol. 22, No.1, 1-34.

FINDLEY, D.F., MONSELL, B.C., BELL, W.R., OTTO, M.C. and CHEN, B.C. (1998), "New Capabilities and Methods of the X12 ARIMA Seasonal Adjustment Program" (with discussion), *Journal of Business and Economic Statistics*, 12, 127-177.

GÓMEZ, V. and MARAVALL, A. (2001a), "Seasonal Adjustment and Signal Extraction in Economic Time Series", Ch.8 in Peña D., Tiao G.C. and Tsay, R.S. (eds.) *A Course in Time Series Analysis*, New York: J. Wiley and Sons.

GÓMEZ, V. and MARAVALL, A. (2001b), "Automatic Modeling Methods for Univariate Series", Ch.7 in Peña D., Tiao G.C. and Tsay, R.S. (eds.), *A Course in Time Series Analysis*, New York: J. Wiley and Sons.

GÓMEZ, V. and MARAVALL, A. (1996), "Programs TRAMO and SEATS. Instructions for the User", (with some updates), Working Paper 9628, Servicio de Estudios, Banco de España.

GÓMEZ, V. and MARAVALL, A. (1994), "Estimation, Prediction and Interpolation for Nonstationary Series with the Kalman Filter", *Journal of the American Statistical Association* 89, 611-624.

GÓMEZ, V., MARAVALL, A. and PEÑA, D. (1999), "Missing Observations in ARIMA Models: Skipping Approach Versus Additive Outlier Approach", *Journal of Econometrics*, 88, 341-364.

HILLMER, S.C. and TIAO, G.C. (1982), "An ARIMA-Model Based Approach to Seasonal Adjustment", *Journal of the American Statistical Association* 77, 63-70.

KAISER, R. and MARAVALL, A. (2005), " Combining Filter Design with Model-based Filtering: An Application to Business-cycle Estimation", *International Journal of Forecasting*, 21, 691-710.

MARAVALL, A. (2003). “A Class of Diagnostics in the ARIMA-model-based Decomposition of a Time Series” in *Seasonal Adjustment*, European Central Bank, November.

MARAVALL, A. (1995), "Unobserved Components in Economic Time Series", in Pesaran, H. and Wickens, M. (eds.), *The Handbook of Applied Econometrics*, chap. 1, 12-72. Oxford: Basil Blackwell.

MARAVALL, A. and PIERCE, D.A. (1987), "A Prototypical Seasonal Adjustment Model", *Journal of Time Series Analysis* 8, 177-193. Reprinted in S. Hylleberg (ed.), Modeling Seasonality, Oxford University Press, 1992.

MARAVALL, A. and PLANAS, C. (1999), "Estimation Error and the Specification of Unobserved Component Models", Journal of Econometrics, 92, 325-353. To be reprinted in P. Newbold and S.J. Leybourne, *Recent Developments in Time Series*, The International Library of Critical Writings in Econometrics, Cheltenham, UK: Edward Elgar Publ.

MARTIN, D.E. and BELL, W.R. (2002), “Computation of Asymmetric Signal Extraction Filters and Mean Squared Error for ARIMA Component Models”, *Statistical Research Division Research Report*, November 2002.

McELROY, T. S. and GAGNON, R. (2006), “Finite Sample Revision Variances for ARIMA Model-Based Signal Extraction”, *Statistical Research Division Research Report*, May 2006.

PIERCE, D.A. (1980), "Data Revisions in Moving Average Seasonal Adjustment Procedures", *Journal of Econometrics* 14, 95-114.

TIAO, G.C. and TSAY, R.S. (1983), "Consistency Properties of Least Squares Estimates of Autoregressive Parameters in ARMA Models", *The Annals of Statistics* 11, 856-871.

TSAY, R.S. and TIAO, G.C. (1984), "Consistent Estimates of Autoregressive Parameters and Extended Sample Autocorrelation Function for Stationary and Non-Stationary Arma Models", *Journal of the American Statistical Association* 79, 84-96.