

TRANSFORMATIONS AND SEASONAL ADJUSTMENT

BY TOMMASO PROIETTI AND MARCO RIANI

Università di Roma “Tor Vergata” and Università di Parma

First Version received June 2007

Abstract. We address the problem of seasonal adjustment of a nonlinear transformation of the original time series, measured on a ratio scale, which aims at enforcing two essential features: additivity and orthogonality of the components. The posterior mean and variance of the seasonally adjusted series admit an analytic finite representation only for particular values of the transformation parameter, e.g. for a fractional Box–Cox transformation parameter. Even if available, the analytical derivation can be tedious and difficult. As an alternative we propose to compute the two conditional moments of the seasonally adjusted series by means of numerical and Monte Carlo integration. The former is both fast and reliable in univariate applications. The latter uses the algorithm known as the ‘simulation smoother’ and it is most useful in multivariate applications. We present two case studies dealing with robust seasonal adjustment under the square root and the fourth root transformation. Our overall conclusion is that robust seasonal adjustment under transformations is feasible from the computational standpoint and that the possibility of transforming the scale ought to be considered as a further option for improving the quality of seasonal adjustment.

Keywords. Structural time series models; Box–Cox transformation; simulation smoother; forward search; numerical integration.

1. INTRODUCTION

The linear Gaussian model plays a central role in statistics; it is well understood and its features depend on the (conditional) first and second moments. Transformations aim at establishing a scale, different from the original measurements, for which the linear Gaussian model holds. For variables measured on a ratio scale with a strictly positive support, Tukey (1957) proposed the power transformation to achieve a model with a simple structure, normal errors and constant error variance; this was later modified by Box and Cox (1964) and incorporated into the model building process, so as to become what is commonly referred to as the ‘Box–Cox transformation’. Since then, transformations have become a key element in regression analysis (see Atkinson, 1985; Cook and Weisberg, 1999). Several modifications have been proposed to deal with negative observations and to extend the support to the transformed observation over the entire real interval (see, among others, John and Draper, 1980; Yeo and Johnson, 2000).

This article deals with the seasonal adjustment of univariate time series under a parametric nonlinear transformation of the original scale that depends on a single

parameter. In particular, we concentrate on the Box–Cox power transformation for positive time series. Nevertheless, our approach is immediately generalizable to other parametric transformations that are continuous and invertible, and to multivariate time series as well. The contribution of this article is twofold: first, we propose a novel implementation of the forward search technique (Atkinson and Riani, 2000) aiming at the robust estimation of the transformation parameter; secondly, conditional on the latter, we provide a general model-based solution to the problem of estimating the seasonally adjusted series in the original scale.

Seasonal adjustment rests upon two basic pillars: *additivity* and *orthogonality* of the seasonal and nonseasonal components. This point is made strongly by Bell and Hillmer (1984, sect. 4.2), who state that ‘someone who does not want to make these assumptions is working on a different problem’. This article focuses on the situation when the two previous requirements are fulfilled on a scale other than the original scale of measurement and provides a model-based solution to the adjustment problem. The seasonal effects are defined in terms of deviations from the underlying level component, so that their average over a yearly span has zero expectation (this feature will be referred to as the *seasonal balance constraint*). Our linear Gaussian workhorse model is an unobserved components model known as the *basic structural model* (Harvey, 1989). This specification is successfully fitted to a wide class of economic time series, the same for which the another very popular model in the analysis of seasonal series, the *airline* model, is appropriate (see Maravall, 1985); furthermore, it is robust to mis-specification of the nonseasonal part of the model. For instance, Riani (1998) shows that, in most cases, the estimate of the seasonal component inside the basic structural model is virtually unaffected by the omission of a stochastic cycle.

Current seasonal adjustment practice does not fully take into account the problem of seasonal adjustment under transformation; only the seasonal adjustment programme of Bell Laboratories (SABL), documented by Cleveland *et al.* (1978), performs the selection of a preliminary power transformation parameter that minimizes the covariance between the level and the seasonal components. The issue of transforming the seasonally adjusted estimates on the original scale is not addressed explicitly. Both X-12-ARIMA (Findley *et al.*, 1998) and Tramo-Seats (Gómez and Maravall, 1997) consider the multiplicative and the additive decompositions of the original series and have in-built procedures to select between them. In the multiplicative case, they produce seasonally adjusted series in the original scale. The X-12-ARIMA program allows for the estimation of an autoregressive integrated moving-average (ARIMA) model with regression effects under the Box–Cox and the logistic transformation, but with the specific purpose of obtaining forecast and backcast extensions by the naïve method, i.e. by simple inversion of the extrapolations made on the transformed scale.

The question has to be raised as to why the transformation problem has not received sufficient recognition in the current seasonal adjustment practice. We can envisage three arguments: the first deals with the seasonal balance constraint, by which the expectation of the sum of the seasonal component over a calendar year is zero. According to a well-established view, the constraint should be enforced on

the original measurement scale; to put it differently, the seasonally adjusted series should have the same expectation (average) as the original series over 12 consecutive monthly observations. This view is held strong by statistical agencies; it is also at the root of the treatment of the problem of seasonal adjustment under transformations by Thomson and Ozaki (2002), who propose *ad hoc* solutions with the specific intent of enforcing the seasonal balance constraint on the original scale.

A second argument deals with contemporaneous aggregation: the seasonally adjusted aggregate should be equal to the aggregated sum of the seasonally adjusted sub-series. The consistency in aggregation requires that the series are not transformed as a necessary (though not sufficient) condition, and thus would not hold for the Box–Cox transformation. A third argument concerns the difficulties and the computational burden linked with the detection of influential observations and/or of the outliers on the transformed scales.

None of these arguments is compelling. Multiplicative adjustment, which is used frequently for economic time series already incorporates a different seasonal balance constraint, which refers to the geometric average, rather than the arithmetic. The view taken in this article is that the stochastic seasonal balance constraint needs to hold only on the transformed scale. The transformation parameter uniquely defines what type of seasonal balance constraint is enforced on the original scale; roughly speaking, if the power transformation parameter is 1, then the balance constraint is additive; if the transformation parameter is equal to 0, it is multiplicative; if the transformation parameter is equal to -1 , the seasonal balance constraint is defined on the reciprocal of the series (which corresponds to the harmonic average). Secondly, the conditions for consistency in cross-sectional aggregation are so stringent that the indirect seasonal adjustment of an aggregate is not very often used in practice. In brief, consistency in aggregation requires that the adjustment is performed additively on the original scale using exactly the same filter. As pointed out by one of the referees, a notable example is US unemployment, which is adjusted indirectly from eight components. As far as the third argument is concerned, in this article we show how it is possible to robustly estimate the transformation parameter and at the same way to evaluate the effect that the different seasons exert on this estimate. In general, prior outlier detection is performed on the original scale before estimating the transformation parameter. It is clear, however, that observations which seem atypical on the original scale may fit completely inside the bulk of the data once the observations have been transformed.

The article is structured as follows: section 2 recalls the basic structural model for the Box–Cox transformed data. Section 3 deals with the robust estimation of the transformation parameter, through the use of the forward search algorithm. The evaluation of the posterior mean and variance of the nonseasonal component is considered in section 4. The availability of closed-form solutions is investigated and approximate solutions reviewed. A more general approach is to evaluate the conditional moments by numerical and Monte Carlo integration using the simulation smoother (de Jong and Shephard, 1995). In section 5 the alternative

estimation methods are applied to a well-known case study concerning the sales of an engineering company (Chatfield and Prothero, 1973), which calls for the fourth root transformation, and the Italian industrial production index for leather and shoes, for which the square root transformation is appropriate. The assessment of the different methods leads to the conclusion that numerical integration is both fast and reliable in univariate applications. We draw our conclusions in section 6.

2. THE BASIC STRUCTURAL MODEL UNDER TRANSFORMATIONS

The parametric linear and Gaussian model that we employ for the adjustment is the basic structural model (BSM henceforth, see Harvey, 1989). The BSM postulates an additive and orthogonal decomposition of a time series into unobserved components representing the trend, seasonality and the irregular component.

We assume that the BSM holds for a transformation $u_t(\lambda)$ of the original time series y_t , depending on a single transformation parameter λ . An important case is the Box–Cox (BC) transformation:

$$u_t(\lambda) = \begin{cases} \frac{y_t^\lambda - 1}{\lambda} & \lambda \neq 0 \\ \ln y_t & \lambda = 0 \end{cases} \quad (1)$$

see Box and Cox (1964). The above transformation is suitable for series measured on a ratio scale, which take only strictly positive values.

The BSM for the transformed series is formulated as follows:

$$u_t(\lambda) = \mu_t + \gamma_t + \sum_{k=1}^K \delta_k x_{kt} + \epsilon_t, \quad t = 1, \dots, T, \quad (2)$$

where μ_t is the trend component, γ_t is the seasonal component, the x_{kt} values are appropriate regressors that account for calendar effects, namely trading days, moving festivals (Easter) and the length of the month, and ϵ_t is a normally and independently distributed irregular component with zero mean and variance σ_ϵ^2 , $\epsilon_t \sim \text{n.i.d.}(0, \sigma_\epsilon^2)$.

The trend component has a local linear representation (Harvey, 1989). The seasonal component can be modelled using a trigonometric representation, such that the seasonal effect at time t arises from the combination of a set of stochastic cycles whose common variance is σ_ω^2 . Alternatively, it is possible to use the so-called Harrison and Stevens (HS) specification which is formulated directly in terms of the effect of a particular season, thereby enhancing flexibility needed to model seasonal heteroscedasticity. For a comparison of the various representations of a seasonal component and a discussion of the implications for forecasting, see Proietti (2000). One of the purposes of this article is to check how the presence of seasonal heteroscedasticity may affect the estimate of the transformation parameter.

Given the value of the transformation parameter (its estimation will be dealt with in the next section), the BSM can be cast in state-space form. The Kalman filter enables the evaluation of the likelihood via the prediction error decomposition. (See Durbin and Koopman, 2001 and Harvey and Proietti, 2005, for a review.) The maximum likelihood estimates can be obtained by a quasi-Newton algorithm, such as the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm (see Press *et al.*, 1992, sect. 10.7).

Finally, conditional on the parameter estimates, the Kalman smoother provides the conditional expectations of the latent components, given all the available observations, along with their conditional variance. These inferences are employed in section 4 to produce estimates of the seasonally adjusted series on the untransformed scale.

3. ROBUST ESTIMATION OF THE TRANSFORMATION PARAMETER AND SEASONAL ADJUSTMENT

As far as the estimation of the transformation parameter is concerned, we can maximise the profile likelihood corrected so as to take into account the change of scale of the observations. Equivalently, we can maximize the uncorrected log-likelihood of the normalized observations $u_t(\lambda) / \prod_t y_t^{\lambda-1}$ (Atkinson, 1985).

An alternative approach, which does not require the computing of the maximum likelihood estimate of λ , is to consider the first-order Taylor series expansion of $u_t(\lambda)$ about a maintained value λ_0 (e.g. 0 or 1): $u_t(\lambda) = u_t(\lambda_0) + (\lambda - \lambda_0)w_t(\lambda_0)$, with $w_t(\lambda_0) = (\partial u_t(\lambda) / \partial \lambda)|_{\lambda=\lambda_0}$.

If for some λ ,

$$u_t(\lambda) = \mu_t + \gamma_t + \sum_k \delta_k x_{kt} + \epsilon_t,$$

then the approximate linear model is

$$u_t(\lambda_0) = \mu_t + \gamma_t + \sum_k \delta_k x_{kt} + \delta^* w_t(\lambda_0) + \epsilon_t, \quad \text{with } \delta^* = \lambda_0 - \lambda.$$

The augmented model is estimated including, among the regressors, the additional variable $w_t(\lambda_0)$. Significant regression denotes the need for a transformation and provides a preliminary estimate of the correct λ as $\hat{\lambda} = \lambda_0 - \hat{\delta}$. The t -test on the additional constructed variable $w_t(\lambda_0)$ is known in the statistical literature as ‘score test statistic for transformation’ (Atkinson, 1985, chap. 6).

Nevertheless, it is well known that the estimated transformation and related test statistic may be sensitive to the presence of one, or several, atypical observations. In addition, it is important to remark that outliers in one transformed scale may not be atypical in another scale. Therefore, it is important to choose a transformation which does not depend on the presence of

particular observations. In this article, in order to provide a robust estimate of the transformation parameter, we use the forward search approach in the way suggested by Atkinson and Riani (2000) and extended to time series by Riani (2004).

The algorithm is both efficient and robust. It is efficient, because it makes use of the Gaussian likelihood machinery underlying the Kalman filter. It is robust, because the outliers enter in the last steps of the procedure and their effect on the statistics of interest can be clearly traced. More generally, this approach allows evaluation of the inferential effect that each time period, either outlying or not, exerts on the fitted model.

One major advantage of the forward search over other high-breakdown techniques is that a number of diagnostic measures can be computed and monitored as the algorithm progresses. In the particular context of transformations, we can monitor the forward plot of the approximate score statistic for testing the significance of the set of constructed variables for different values λ_0 , using a separate search for each λ_0 . The trajectories of the score tests can be combined in a single picture named the ‘fan plot’ (Atkinson and Riani, 2002). If the number of observations is not large (i.e. less than 200), generally the five most common values of λ_0 ($-1, -0.5, 0, 0.5, 1$) are sufficient for selecting the appropriate transformation. On the other hand, when the sample size is large we have to consider a finer grid of values of λ_0 . The monitoring of the fan plot for the different specifications of the seasonal component (trigonometric HS or heteroscedastic HS) inside the basic structural model enables to appraise how robust is our estimate of the transformation parameter to the various parameterizations of γ_t .

An additional novelty of this article is that we implement for the first time in time series the so-called ‘proportional forward search’ in order to achieve balance in observations across seasons during the search. More precisely, at each step, among observations for the seasons least represented in the current set, we add the observation with the smallest one-step-ahead standardized prediction residual. In this way in each step the subset has a composition of months, which reflects as much as possible the structure of the overall sample.

4. SEASONAL ADJUSTMENT AND THE BOX-COX TRANSFORMATION

Let us write eqn (2) as

$$u_t(\lambda) = u_t^* + \gamma_t + \sum_k \delta_k x_{kt},$$

where $u_t^* = \mu_t + \epsilon_t$ is the seasonally adjusted series on the transformed scale, and denote by $\tilde{u}_t^* = E(u_t^* | \mathcal{F}_T)$ and $V_t = \text{var}(u_t^* | \mathcal{F}_T)$ the posterior mean and variance of u_t^* , respectively, \mathcal{F}_T being the information set at time t . These inferences are delivered by the Kalman filter and smoother applied to the relevant linear state space model. (See, e.g. Durbin and Koopman, 2001, for details.)

We define the seasonally adjusted series on the original scale as the inverse transformation of the nonseasonal component $u_t^*, y_t^* = u^{-1}(u_t^*)$, where $u^{-1}(\cdot)$ is the inverse transformation. For the Box–Cox transformation:

$$y_t^* = \begin{cases} (1 + \lambda u_t^*)^{1/\lambda}, & \lambda \neq 0, \\ \exp(u_t^*), & \lambda = 0. \end{cases}$$

The estimator of the seasonally adjusted series is thus:

$$\tilde{y}_t^* = E(y_t^* | \mathcal{F}_T) = \int u^{-1}(u_t^*) f(u_t^* | \mathcal{F}_T) du_t^*, \tag{3}$$

whereas the conditional variance of the estimation error for the seasonally adjusted series is defined as:

$$\text{var}(y_t^* | \mathcal{F}_T) = \int [u^{-1}(u_t^*) - \tilde{y}_t^*]^2 f(u_t^* | \mathcal{F}_T) du_t^* = E(y_t^{*2} | \mathcal{F}_T) - \tilde{y}_t^{*2}. \tag{4}$$

The above integrals do have a closed-form solution only in particular cases, namely $\lambda = 0$, and $\lambda = 1/p$, $p = 1, 2, 3, \dots$, as will be seen shortly.

Notice that the naïve estimator of the SA series,

$$\hat{y}_t^* = \begin{cases} (1 + \lambda \tilde{u}_t^*)^{1/\lambda}, & \lambda \neq 0, \\ \exp(\tilde{u}_t^*), & \lambda = 0, \end{cases} \tag{5}$$

provides the *median* of the conditional distribution of y_t^* , given the observations.

The naïve estimator has an interpretative advantage over eqn (3) in the case $\lambda = 0$, as y_t can be decomposed exactly as the product of $\hat{y}_t = \exp(\tilde{u}_t^*)$ and the exponential of the estimate of the seasonal component, $\exp[E(\gamma_t + \sum_k \delta_k x_{kt} | \mathcal{F}_T)]$. Except for the linear-additive case, the minimum mean square estimator of the seasonally adjusted series (3) cannot be combined with the minimum mean square estimator of the seasonal component in the original scale to provide an exact decomposition of the untransformed time series, y_t . Thomson and Ozaki (2002) propose *ad hoc* adjustments to the estimated components that aim at enforcing the seasonal balance constraint in the original scale; their proposal may be adapted to the estimators of the components proposed in this article.

4.1. Analytical solutions

For general λ , in the Appendix we prove Theorem 1.

THEOREM 1. *The mean and the variance of the seasonally adjusted series in the original scale are given by the two following expressions:*

$$E(y_t^* | \mathcal{F}_T) = \tilde{y}_t^* = \hat{y}_t^* \left[1 + \sum_{j=1}^{\infty} k_{2j}(t) a_j(t) \right] \tag{6}$$

$$\text{var}(y_t^* | \mathcal{F}_T) = \hat{y}_t^{*2} \left[\sum_{j=1}^{\infty} k_j^2(t) a_j(t) + 2 \sum_{j=1}^{\infty} \sum_{r=1}^{\infty} k_j(t) k_{j+2r}(t) a_{j+r}(t) - \left(\sum_{j=1}^{\infty} k_{2j}(t) a_j(t) \right)^2 \right] \tag{7}$$

where¹

$$a_j(t) = E \left[(u_t^* - \tilde{u}_t^*)^{2j} | \mathcal{F}_T \right] = \frac{(2j)!}{j! 2^j} V_t^j \quad \text{and} \quad k_j(t) = \frac{1}{j!} \left(\prod_{k=1}^{j-1} (1 - \lambda k) \right) \hat{y}_t^{*-\lambda j}$$

The results follow from the Taylor series expansion of the reverse transformation. Notice that for $\lambda = 0$ $k_j(t) = (j!)^{-1}$ and the term of eqn (6) is simply the expansion of $\exp(V_t/2)$. This method was proposed originally by Neyman and Scott (1960), who however did not consider explicitly time series applications and did not give the exact analytical solution for $\lambda = 1/p$, with p an integer. An alternative approach for expressing the time-series forecasts on the original scale, based on Hermite polynomial expansion, was suggested by Granger and Newbold (1976). The expression (6) was derived by Pankratz and Dudley (1987) for the simple power transformation y_t^{λ} using a different argument. For integer $p = 1/\lambda$ they write the inverse transformation as

$$u_t^{*p} = (\tilde{u}_t^* + \sqrt{V_t} w_t)^p = \tilde{u}_t^{*p} \left(1 + \frac{\sqrt{V_t}}{\tilde{u}_t^*} w_t \right)^p, \quad \text{where } w_t \sim N(0, 1).$$

They then consider the expansion of the binomial and take the expectation. The expressions in square brackets in eqns (6) and (7) are the multiplicative correction terms that have to be applied to the naïve estimator of the SA series or to its square in order to produce the conditional mean and the conditional variance in the original scale.

An alternative expression for the variance is derived as follows. Defining \hat{V}_t^* as the naïve as estimate of the variance resulting from the application of the delta method,

$$\hat{V}_t^* = V_t \left[\frac{du^{-1}(u_t^*)}{du_t^*} \Big|_{u_t^* = \tilde{u}_t^*} \right]^2 = V_t \hat{y}_t^{*2(1-\lambda)},$$

then we can rewrite eqn (7) as:

$$\text{var}(y_t^* | \mathcal{F}_T) = \hat{V}_t^* \left[1 + \sum_{j=2}^{\infty} \bar{k}_j^2(t) \bar{a}_j(t) + 2 \sum_{j=1}^{\infty} \sum_{r=1}^{\infty} \bar{k}_j(t) \bar{k}_{j+2r}(t) \bar{a}_{j+r}(t) - V_t \left(\sum_{j=1}^{\infty} \bar{k}_{2j}(t) \bar{a}_j(t) \right)^2 \right], \tag{8}$$

where $\bar{k}_j(t) = k_j(t) \hat{y}_t^{*\lambda}$ and $\bar{a}_j(t) = a_j(t) / V_t$. According to expression (8), the exact variance can be seen as the product of the naïve variance resulting from the delta method and a correction factor.

For $\lambda = 1/p, p = 1, 2, \dots$, we can immediately see that the series $k_1(t), k_2(t), \dots$ contains only p terms different from zero. We give in Table I, for the most common values of λ , the exact correction factors for the mean and the variance which must be applied to the naïve estimator of the seasonally adjusted series \hat{y}_t^* in order to find the true conditional mean and variance in the original scale. This table clearly shows that the correction term depends on the ratio between the variance (raised to some power) of the SA series on the transformed scale and the value of the naïve estimator (raised to some power of λ). If this is small, the correction is negligible. More precisely, we have Corollary 1.

COROLLARY 1. *The correction factor for the mean which we call $\psi_\mu(\lambda, \hat{y}_t^*, V_t)$ satisfies the following properties:*

- (i) $\psi_\mu(\lambda, \hat{y}_t^*, V_t) \leq 1$ for $\lambda \geq 1$ and $\psi_\mu(\lambda, \hat{y}_t^*, V_t) \geq 1$ for $\lambda \leq 1$.
- (ii) $\psi_\mu(\lambda, \hat{y}_t^*, V_t) \rightarrow 1^-$ when $\lambda \rightarrow +\infty$ and $\psi_\mu(\lambda, \hat{y}_t^*, V_t) \rightarrow +\infty$ when $\lambda \rightarrow -\infty$ if $\hat{y}_t^* > 1$. $\psi_\mu(\lambda, \hat{y}_t^*, V_t) \rightarrow -\infty$ when $\lambda \rightarrow +\infty$ and $\psi_\mu(\lambda, \hat{y}_t^*, V_t) \rightarrow 1^+$ when $\lambda \rightarrow -\infty$ if $\hat{y}_t^* < 1$.

The proofs are straightforward but tedious. Additional details and other properties are given in a technical report by the authors (Proietti and Riani, 2007).

Figure 1 shows the correction factor as a function of λ for six different combinations of values of \hat{y}_t^* and V_t . It is clear that if the correction factor is neglected there is negative (positive) bias for $\lambda < 1 (\lambda > 1)$ which can be more or less severe depending on the problem under study. The first two left panels show that if the ratio between the naïve estimator and the value of the variance of the SA series in the transformed scale is greater than a certain threshold and if the estimated λ is greater (smaller) than 1 and \hat{y}_t^* is greater (smaller) than 1, the correction which must be applied to the naïve estimator can be overlooked.

TABLE I

EXACT CORRECTION FACTORS WHICH HAVE TO BE APPLIED TO THE NAIÏVE ESTIMATOR OF THE SEASONALLY ADJUSTED SERIES AND OF THE VARIANCE, IN ORDER TO OBTAIN THE CONDITIONAL MEAN AND THE CONDITIONAL VARIANCE IN THE ORIGINAL SCALE FOR THE MOST IMPORTANT FRACTIONAL VALUES OF λ AND FOR $\lambda = 0$

	Mean	Variance
λ	Correction factor for \hat{y}_t^*	Correction factor for \hat{V}_t^*
1/2	$1 + \frac{1}{4} \frac{V_t}{\hat{y}_t^*}$	$1 + \frac{1}{8} V_t \hat{y}_t^{*-1}$
1/3	$1 + \frac{1}{3} \frac{V_t}{\hat{y}_t^{2/3}}$	$1 + \frac{4}{9} V_t \hat{y}_t^{*-2/3} + \frac{5}{243} V_t^2 \hat{y}_t^{*-4/3}$
1/4	$1 + \frac{3}{8} \frac{V_t}{\hat{y}_t^{1/2}} + \frac{3}{256} \frac{V_t^2}{\hat{y}_t}$	$1 + \frac{21}{32} V_t \hat{y}_t^{*-1/2} + \frac{3}{32} V_t^2 \hat{y}_t^{*-1} + \frac{3}{32} V_t^3 \hat{y}_t^{*-3/2}$
0	$\exp(\frac{V_t}{2})$	$V_t^{-1} \exp(V_t) \cdot (\exp(V_t) - 1)$

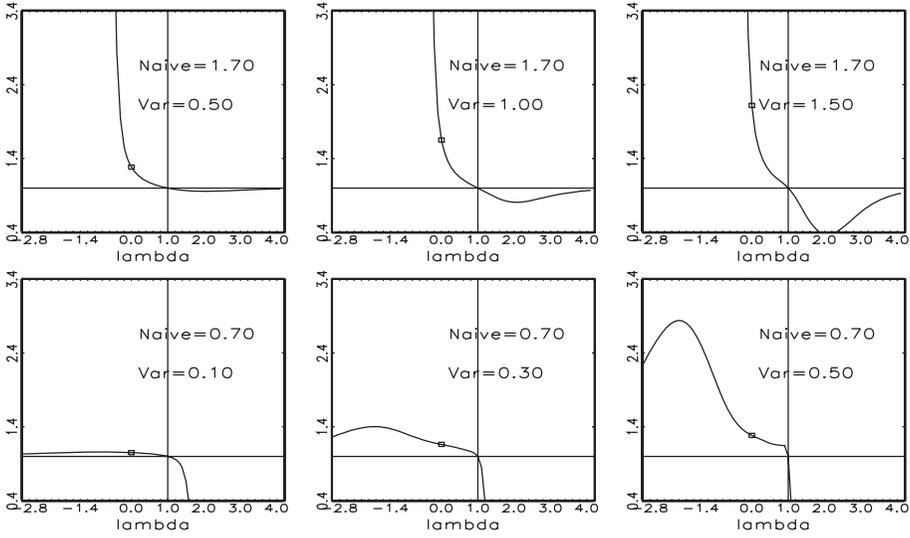


FIGURE 1. Correction factor which must be applied to the naïve estimator of the seasonally adjusted series to obtain the conditional mean in the original scale as a function of λ for six different combinations of \hat{y}_t^* (naïve) and V_t (VAR). The square which is drawn in correspondence of $\lambda = 0$ denotes the value obtained by directly applying the formula given in Table I.

LEMMA 2. *The correction factor for the naïve variance which we call $\psi_\sigma(\lambda, \hat{y}_t^*, V_t)$ satisfies the following property:*

- (i) $\psi_\sigma(\lambda, \hat{y}_t^*, V_t) \rightarrow 1^+$ when $\lambda \rightarrow +\infty$ and $\psi_\sigma(\lambda, \hat{y}_t^*, V_t) \rightarrow +\infty$ when $\lambda \rightarrow -\infty$ if $\hat{y}_t^* > 1$. $\psi_\sigma(\lambda, \hat{y}_t^*, V_t) \rightarrow 1^+$ when $\lambda \rightarrow -\infty$ and $\psi_\sigma(\lambda, \hat{y}_t^*, V_t) \rightarrow +\infty$ when $\lambda \rightarrow +\infty$ if $\hat{y}_t^* < 1$.

which is illustrated graphically in Figure 2.

4.2. Approximate and computational solutions

Taylor (1986) proposed an approximate correction for the case $\lambda \neq 0$, which amounts to neglecting higher order terms in the expansion (6):

$$\tilde{y}_t^{*T} = \hat{y}_t^* \left[1 + \frac{1}{2}(1 - \lambda) \frac{V_t}{\hat{y}_t^{*2\lambda}} \right].$$

This estimate is exact only for $\lambda = 0.5$ (see Table I).

Guerrero (1993) proposed a solution which is coincident with the exact solution in the logarithmic case ($\lambda = 0$) and is approximate for $\lambda \neq 0$. In our notation, it can be written as follows:

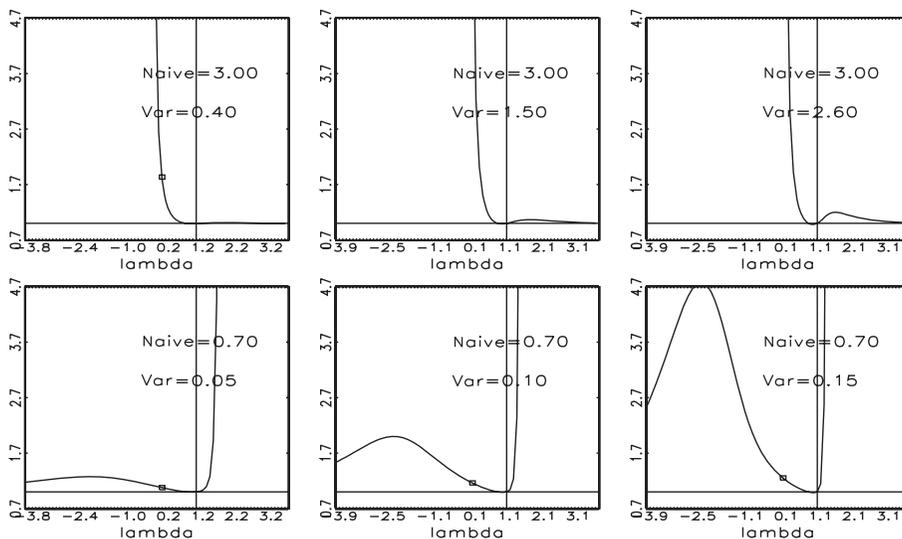


FIGURE 2. Correction factor which must be applied to the naïve estimator of the variance of the seasonally adjusted series to obtain the conditional variance in the original scale as a function of λ for six different combinations of \hat{y}_t^* (naïve) and V_t (VAR). The symbol of the square which is drawn in correspondence of $\lambda = 0$ denotes the value obtained applying directly the formula given in Table I.

$$\tilde{y}_t^{*G} = \tilde{y}_t^* \left\{ \frac{1}{2} + \frac{1}{2} \left[1 + 2\lambda(1 - \lambda) \frac{V_t}{\hat{y}_t^{*2\lambda}} \right]^{1/2} \right\}^{1/\lambda}.$$

For general λ there are three possible ways of evaluating $E(y_t^*|\mathcal{F}_T)$ and $\text{var}(y_t^*|\mathcal{F}_T)$:

- Monte Carlo evaluation using the simulation smoother: the latter is used to draw repeated samples from the conditional distribution of $u^* = \{u_1^*, \dots, u_T^*\}$, given the available observations.
- Numerical integration with respect to the normal density, $f(u_t^*|\mathcal{F}_T)$, whose moments \tilde{u}_t^* and V_t^* are provided by the Kalman filter and smoother.
- Direct application of eqns (6) and (7) truncating the summations to a particular order.

For the Monte Carlo evaluation, we use the method of de Jong and Shephard (1995). It achieves computational efficiency by sampling from the joint posterior density of the disturbances in the model. An even more efficient method is the simulation smoother proposed by Durbin and Koopman (2002).

5. ILLUSTRATIONS

In this section we propose two illustrations dealing with seasonal adjustment under the square root transformation and $\lambda = 1/4$. All the computations were

performed using Ox 3.x by Doornik (2001) and the library of state-space function SSFPAK 2.3 by Koopman *et al.* (1999). The numerical integration for eqn (8) is implemented using the QUADPACK function QAGS (see Piessens *et al.*, 1983; QUADPACK is a Fortran library for univariate numerical integration ‘quadrature’ using adaptive rules).

5.1. Sales X data

Our first illustration deals with a well-known case study, concerning the monthly sales of a engineering company (company X), from January 1965 to May 1971, that was presented and studied by Chatfield and Prothero (1973) as a case study by using Box–Jenkins forecasting methods. The plot of the series (see the first panel of Figure 3) reveals that the amplitude of the seasonal pattern is increasing over time as the trend increases, but the evidence is that the logarithmic transformation is overtransforming the series, i.e. the amplitude decreases as the trend increases on the transformed scale.

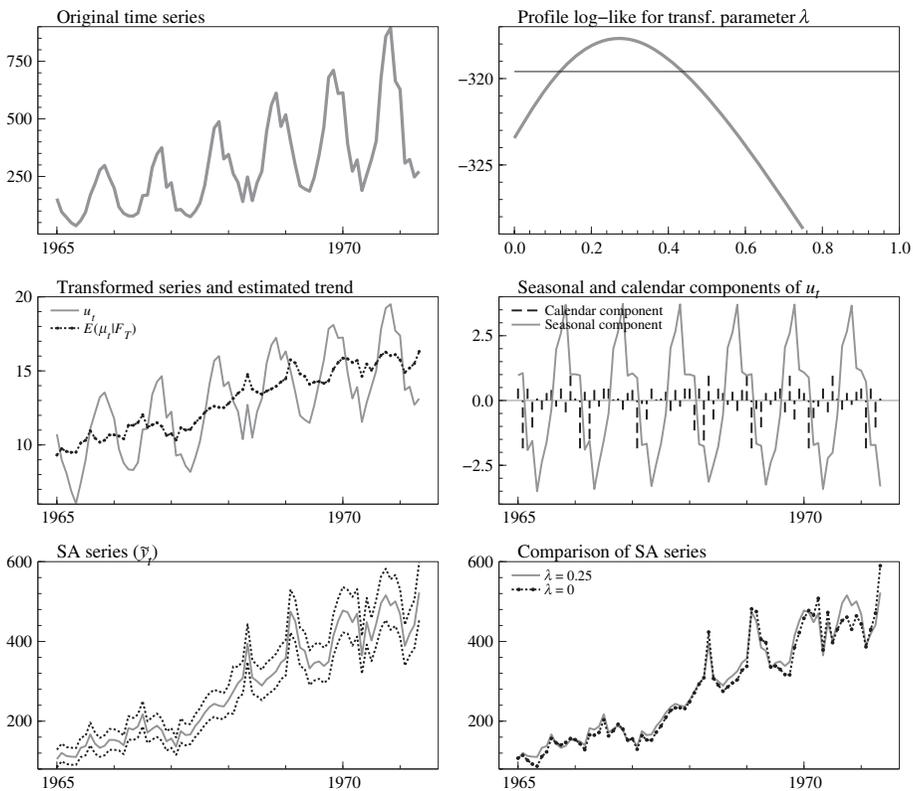


FIGURE 3. Sales of Company X.

In their discussion of the Chatfield and Prothero paper, Box and Jenkins (1973) using range-mean plots suggested the value 0.25, which is also the value estimated by Chatfield and Prothero in their reply. That value is confirmed by Guerrero (1993), by a different method, which looks at the variance-stabilizing properties of the transformation.

When the BSM is estimated under the Box–Cox transformation, the profile likelihood for the parameter λ is reported in the second panel of Figure 3. The horizontal line is drawn at $\mathcal{L}_{\max} - 0.5\chi^2(0.95)$, where \mathcal{L}_{\max} is the value of the corrected profile likelihood evaluated at the maximum and $\chi^2(0.95)$ is the 95th percentile of the χ^2 distribution with 1 degree of freedom (3.84).

The logarithmic transformation and the value $\lambda = 1$ are clearly rejected and the maximum likelihood estimate is $\tilde{\lambda} = 0.27$. It is worth noticing that the differences in results with respect to other authors can be attributed to the fact that we use a different model and that we include a calendar component in our model, which turns out to be significant. As the value 0.27 is not significantly different from 0.25, in our subsequent analyses we will use the value $\lambda = 0.25$, for which, as we have seen in the previous section, the conditional mean and the variance of the seasonally adjusted series admit a closed-form solution.

The maximum likelihood estimates of the variance parameters are $\tilde{\sigma}_\eta^2 = 0.1108$; $\tilde{\sigma}_\zeta^2 = \tilde{\sigma}_\omega^2 = 0.0000$; and $\tilde{\sigma}_\epsilon^2 = 0.1728$. As a result, the slope is fixed and seasonality is deterministic. The Bowman and Shenton normality test takes the value 0.5; some residual correlation is left, as by the Ljung–Box portmanteau test statistic with 12 autocorrelation, which takes the value 23.34.

The two central panels present the transformed series along with the estimated trend and the seasonal and calendar component on the transformed scale. The bottom panel displays the estimates of the seasonally adjusted series on the original scale, that is \tilde{y}_t^* along with their 95% highest density region. The computations were made by numerical integration, but as we argue below these are undistinguishable from the exact estimates and from the Monte Carlo estimates using a suitably large number of replications. It is interesting to notice, as we have seen theoretically in the previous section, that the width of the confidence interval of the seasonally adjusted series in the original scale increases as the trend increases.

The last panel compares the estimates of the SA series arising for the estimated transformation parameter with that arising in the case of the logarithmic transformation. The graph highlights that the differences can be relevant and the Box–Cox transformation is indeed an issue in seasonal adjustment. Given that an exact solution exists for \tilde{y}_t^* and $\text{var}(y_t^*|\mathcal{F}_T)$, we can evaluate the accuracy of the various estimates that have been proposed. The estimation methods that are compared are the following.

- The naïve estimate (the conditional median), $\tilde{y}_t^* = (1 + 0.25\tilde{u}_t^*)^4$.
- Monte Carlo integration using the simulation smoother: M independent samples, $u_t^{(i)*}$, $i = 1, \dots, M$, are drawn from the conditional distribution

$u_t^* | \mathcal{F}_T \sim N(\tilde{u}_t^*, V_t^*)$, which is done recursively by the simulation smoother. The seasonally adjusted series is estimated by average

$$\tilde{y}_t^{*MC} = \frac{1}{M} \sum_{i=1}^M \left[1 + 0.25u_t^{(i)*} \right]^4.$$

The variance of the SA series is estimated by

$$\text{v\`ar}(y_t^* | \mathcal{F}_T) = \frac{1}{M} \sum_{i=1}^M \left[1 + 0.25u_t^{(i)*} \right]^8 - (\tilde{y}_t^{*MC})^2.$$

Results are presented for the number of replications $M = 1000$ and $M = 10,000$.

- Numerical integration using the QUADPACK function QAGS, available in Ox 3.4; the finite integration interval is defined as $[\tilde{u}_t^* - 8\sqrt{V_t}, \tilde{u}_t^* + 8\sqrt{V_t}]$, where \tilde{u}_t and V_t are evaluated by the Kalman filter and smoother applied to the transformed observations.
- The Taylor estimation method based on a Taylor’s approximation:

$$\tilde{y}_t^{*T} = \hat{y}_t^* \left[1 + \frac{3}{8} \frac{V_t}{\sqrt{\hat{y}_t^*}} \right].$$

- The method proposed by Guerrero:

$$\tilde{y}_t^{*G} = \hat{y}_t^* \left\{ \frac{1}{2} + \frac{1}{2} \left[1 + \frac{3}{8} \frac{V_t}{\sqrt{\hat{y}_t^*}} \right]^{1/2} \right\}^4.$$

Based on the choice of model and λ , Table II reports the mean error of method j ,

$$\text{ME}_j = T^{-1} \sum_{t=1}^T (\tilde{y}_t^{*j} - \tilde{y}_t^*),$$

TABLE II
SALES X DATA: ACCURACY OF DIFFERENT ESTIMATION METHODS

Method	Mean error	Mean square error	Mean percent error	MAPE
Estimation of $\tilde{y}_t^* = E(y_t^* \mathcal{F}_T)$				
Naive	-0.48452941	0.26217669	-0.18820985	0.18820985
MC Int 1000	-0.00587377	0.00044615	-0.00227596	0.00646891
MC Int 10000	0.00173609	0.00003516	0.00060559	0.00175783
Num Int	0.00000000	0.00000000	0.00000000	0.00000000
Taylor	-0.00007648	0.00000001	-0.00003273	0.00003273
Guerrero	0.00003824	0.00000000	0.00001636	0.00001636
Estimation of $\text{Var}(y_t^* \mathcal{F}_T)$				
MC Int 1000	10.1550	440.8200	1.9854	3.5628
MC Int 10000	-1.1800	34.7828	-0.4323	0.9942
Num Int	0.0000	0.0000	0.0000	0.0000

where the subtrahend is given by the exact expression given in Table I, the mean square error,

$$\text{MSE}_j = T^{-1} \sum_{t=1}^T (\tilde{y}_t^{*j} - \tilde{y}_t^*)^2,$$

the mean percent error,

$$\text{MPE}_j = 100T^{-1} \sum_{t=1}^T [(\tilde{y}_t^{*j} - \tilde{y}_t^*)/\tilde{y}_t^*]$$

and the mean absolute percent error,

$$\text{MAPE}_j = 100T^{-1} \sum_{t=1}^T [|\tilde{y}_t^{*j} - \tilde{y}_t^*|/\tilde{y}_t^*].$$

In this application, the ratio $V_t/\sqrt{\hat{y}_t^*}$ is very small (6×10^{-5} on average) and thus the naïve estimate has a good performance. It should be recalled that the last two columns present percent values. It is also evident from the table that the Taylor and Guerrero approximations are very accurate for this application. Numerical integration is the most accurate; the performance of Monte Carlo integration depends on the number of replications used. The convergence to the true conditional mean is not very fast. This is due to the correlation between the random draws that results from the persistence of the nonseasonal component of the series. The use of an antithetic variable greatly improves the performance.

The second part of the table displays the same statistics with reference to the problem of estimating the conditional variance $\text{var}(y_t^*|\mathcal{F}_T)$. It must be remarked that the Taylor and Guerrero methods do not provide an estimate of this feature. Again, numerical integration provides the fastest and most reliable method of estimating $\text{var}(y_t^*|\mathcal{F}_T)$.

5.2. Italian industrial production of LS sector

Our second illustration deals with the estimation of the seasonally adjusted series and of its posterior variance according to eqns (3) and (4) with reference to the industrial production index for the Leather and Shoes (LS) sector, available for the period 1981.1–2005.2 (source Istat, base 2000 = 100, 290 observations), under the Box–Cox transformation. We notice in passing that the official seasonal adjustment performed by the Italian National Statistical Institute (Istat) is carried out on the untransformed series (i.e. $\lambda = 1$) using the software Tramo-Seats.

The plot of the original series (see the left-hand panel of Figure 4) reveals that the amplitude of the seasonal component decreases with the trend. The dominant feature is the seasonal trough occurring in August. The likelihood ratio test of $H_0:\lambda = 1$ is significant and the maximum likelihood estimate of the transformation parameter is $\hat{\lambda} = 0.501$, corresponding to the square root

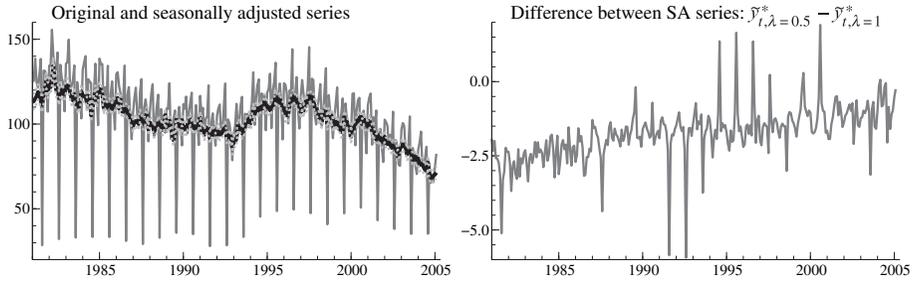


FIGURE 4. Index of Industrial Production, Sector DC: Leather and Shoes.

transformation. The profile likelihood for the transformation parameter suggests values roughly between 0.25 and 0.75.

Clearly, we have to establish whether the square root transformation is due to the presence of particular observations or it is diffused throughout the data. Finally, we need to know what is the effect on the estimated λ of the months of August or whether there are other months whose variance of the seasonal movements is much greater than the others, but are obscured by the high fluctuations of the month of August.

To start answering all these questions in Figure 5 we produce a series of fan plots for $\lambda = (0,0.25,0.5,0.75)'$. The top left panel of Figure 5, which uses a

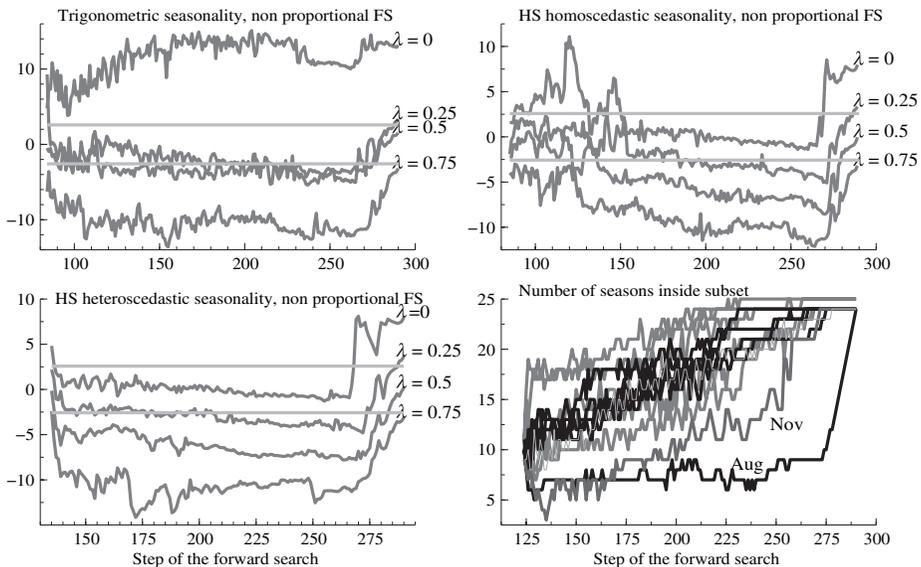


FIGURE 5. Robust estimate of the transformation parameter for different specifications of the seasonal component using a nonproportional FS.

trigonometric specification for the seasonal component and a nonproportional forward search, shows that the log transformation is always rejected throughout. The value 0.75 is strongly rejected during the search until a set of observations at the end brings the score close to the acceptance region. The same upward trend is visible in the curves associated with $\lambda = 0.25$ and $\lambda = 0.5$. In these cases, this set of observations brings the values of the score from value -3 to a value around 3.

The monitoring of the seasons inside subset (bottom right panel) clearly shows that the units entering the subset in steps 270–290 all belong to the month of August. The effect of the month of August is even more pronounced if we consider the HS specification for stochastic seasonality (see top right panel in Figure 5).

In order to understand whether this is due to seasonal heteroscedasticity we redo the fan plot allowing the variance of the month of August to be different from that of the other months. The resulting fan plot, which is given in the bottom left panel, shows that the presence of heteroscedastic seasonality for the month of August does not seem to alter our conclusions about the transformation parameter.

A major benefit of the fan plot is that it clearly enables us to appreciate the effect that the different months and/or different subperiods exert on the estimate of the transformation parameter. As is well known, the FS provides an ordering of the data from those most in agreement with a suggested model (which enter the first steps) to those least in agreement with it (which are included in the final steps). For example, the bottom right panel shows that the seasons which are most difficult to model are those associated with the months of November and August. However, while the effect of the introduction of the months of November (steps 230–260) does not change the value of the score test appreciably, it is clear as regards the effect that the month of August exerts on the estimated λ .

Figure 6 shows the new fan plot for trigonometric (top panel) and HS specification (bottom panel) for a proportional forward search. Both plots show that if we consider subsets which contain the same proportion of months as that of the original sample, the curves for the different values of λ are more stable and the values associated with the square root transformation in the central and final part of the search always lie inside the confidence bands.

The monitoring of the estimates of the hyperparameters on the square root scale (not given here for lack of space) show that in this scale the values of the variances of the underlying components are stable together with the t -statistics for the trading days and there are no sudden jumps because of the presence of atypical observations.

As a result of this analysis, the BSM was used on the transformed observations $u_t = 2(y_t^{1/2} - 1)$. The use of the trigonometric seasonal specification gave the following maximum likelihood estimates of the variance parameters $\tilde{\sigma}_\eta^2 = 0.01556$; $\tilde{\sigma}_\zeta^2 = 0.00003$; $\tilde{\sigma}_\omega^2 = 0.00079$; and $\tilde{\sigma}_\epsilon^2 = 0.05640$. There is a significant calendar component in the series, the coefficients associated with the working days being positive and those associated with the week-end being negative. The diagnostics are satisfactory, and normality is

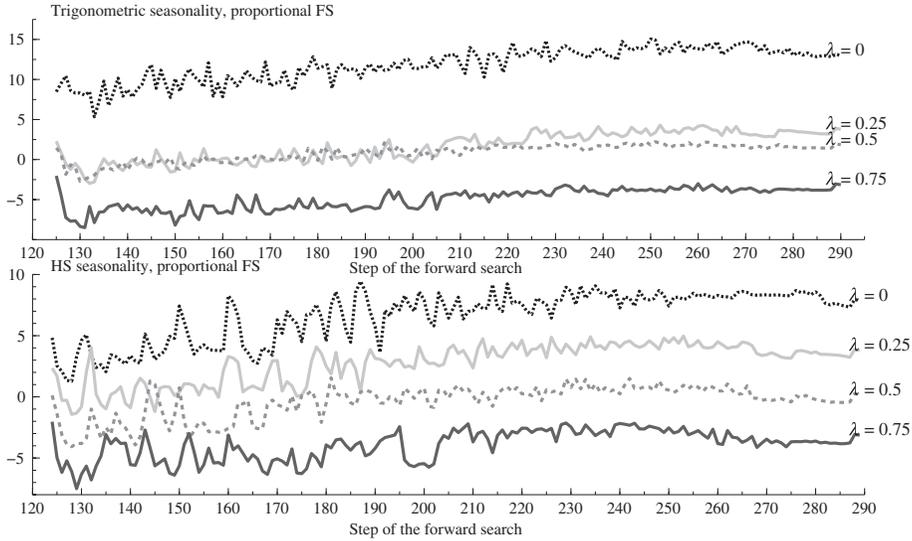


FIGURE 6. Robust estimate of the transformation parameter for different specifications of the seasonal component using a proportional FS.

accepted (the Bowman and Shenton normality test results 3.76, with p -value = 0.15). The point estimates of the seasonally adjusted series on the original scale of measurement are reproduced in the bottom left panel of Figure 4, along with the 95% highest density region of the posterior distribution of y_t^* . The right-hand panel of Figure 4 compares the estimates of the seasonally adjusted series on the original scale with those emerging from the BSM adapted to the untransformed series ($\lambda = 1$); on average, the latter display a positive bias, but there are relevant differences that go well beyond a level change. In particular, the differences are substantial with respect to August.

Given that analytical solutions are available,

$$\tilde{y}_t^* = \hat{y}_t^* \left[1 + \frac{1}{4} \frac{V_t}{\hat{y}_t^*} \right]$$

and

$$\text{var}(y_t^* | \mathcal{F}_T) = \hat{y}_t^* V_t + \frac{1}{8} V_t^2,$$

we can assess the accuracy of the various estimation methods considered in section 5.1. It must be stressed that the Taylor method gives only in this case ($\lambda = 0.5$) an exact solution. The ME, the MSE, the MPE and the MAPE for the seasonally adjusted series and its variance (not reported for brevity), show that the most accurate method is numerical integration, which has an excellent performance also for the estimation of the conditional variance. Monte Carlo integration is more accurate than Guerrero’s method and both outperform the naïve estimate.

6. CONCLUDING REMARKS

This article has investigated the issue of seasonal adjustment under the Box–Cox power transformation of time series which are bounded from below by 0. For lack of space we could only show example from univariate time series. The reader interested in the extension of the current approach to multivariate time series and or to other classes of transformation, such as the Aranda-Ordaz, can refer to a research report by the same authors (Proietti and Riani, 2007).

The rationale behind the transformation is to enhance several desirable features of the maintained measurement model: linearity, additivity and orthogonality of components, normality of the disturbances driving the components. In this article we have concentrated on the Box–Cox transformation applied to the basic structural model. However, the idea of imposing the seasonal constraint on the transformed scale, performing seasonal adjustment and then transforming back into the original series can be applied to any structural or ARIMA model or to the more complicated models like the so-called ‘transformation/weighting’ (see, e.g. Carroll and Ruppert, 1991) models, where not only the response is transformed, but also the part used to fit the mean model and the disturbance term to take into account heteroscedasticity.

This article has documented that transformations of seasonal time series are both feasible and relevant. It is feasible, since there are computationally efficient and accurate methods of estimating the conditional mean and variance of the seasonally adjusted series that are applicable in the absence of a closed-form solution. It is relevant, since the estimates may differ relevantly from those obtained using either the untransformed observations or the logarithms. Our case studies concerned cases when seasonality is the most prominent source of variation of the data as often occurs for industrial production, tourism and sales.

For univariate analysis numerical integration is both fast and reliable, and is recommended against the use of approximate methods or Monte Carlo integration using the simulation smoother. The latter may require a large number of replications, when the nonseasonal component is highly persistent and weakly evolutive. However, it can be made as accurate as needed by increasing the number of replications and by using variance reduction techniques.

While in univariate models it is possible to choose between different solutions, in multivariate applications the use of the simulation smoother is the unique option available. A multivariate application that we have in mind is indirect seasonal adjustment, when a cross sectional seasonally adjusted aggregate is obtained as

$$Y_t^* = \sum_{i=1}^N (1 + \lambda_i u_{it}^*)^{1/\lambda_i}.$$

In this setting, $E(Y_t^* | \mathcal{F}_T)$ can be evaluated by Monte Carlo integration, using the simulation smoother. Applying different transformations to the component series could raise questions about the implied seasonal component; this may

suggest the restriction $\lambda_i = \lambda$, where λ is the transformation parameter of the aggregate series.

Finally, it is worthwhile to remark that even if the focus of this paper was seasonal adjustment, our method can be easily extended to find the estimate of all the other components on the original scale (e.g. the detrended series). In other words, once the two conditional moments of the detrended series in the transformed scaled are found using the Kalman filter and smoother, the detrended series on the original scale can be computed using numerical or Monte Carlo integration or the exact analytic solution described in the article.

APPENDIX: PROOF OF THEOREM 1

We start considering the Taylor series expansion of the inverse transformation $(1 + \lambda u_t^*)^{1/\lambda}$ around \tilde{u}_t^* :

$$\begin{aligned} (1 + \lambda u_t^*)^{1/\lambda} &= (1 + \lambda \tilde{u}_t^*)^{1/\lambda} + (1 + \lambda \tilde{u}_t^*)^{1/\lambda-1} (u_t^* - \tilde{u}_t^*) \\ &\quad + \frac{1}{2} (1 - \lambda) (1 + \lambda \tilde{u}_t^*)^{1/\lambda-2} (u_t^* - \tilde{u}_t^*)^2 + \dots \\ &= \hat{y}_t^* \left[1 + \sum_{j=1}^{\infty} \frac{1}{j!} \left(\prod_{k=1}^{j-1} (1 - \lambda k) \right) (\hat{y}_t^*)^{-\lambda j} (u_t^* - \tilde{u}_t^*)^j \right] \end{aligned}$$

Now, taking the expectation of both sides with respect to the Gaussian density $f(u_t^* | \mathcal{F}_T)$ and remembering that the central j -order moment is zero if j is odd, after some manipulation we obtain that:

$$E(y_t^* | \mathcal{F}_T) = \tilde{y}_t^* = \hat{y}_t^* \left[1 + \sum_{j=1}^{\infty} \frac{1}{j! 2^j} \left(\prod_{k=1}^{2j-1} (1 - \lambda k) \right) \frac{V_t^j}{\hat{y}_t^{*2\lambda j}} \right]. \tag{9}$$

If we denote with

$$a_j(t) = \frac{(2j)!}{j! 2^j} V_t^j \quad \text{and} \quad k_j(t) = \frac{1}{j!} \left(\prod_{k=1}^{j-1} (1 - \lambda k) \right) \hat{y}_t^{*-\lambda j},$$

eqn (9) can be rewritten as:

$$E(y_t^* | \mathcal{F}_T) = \tilde{y}_t^* = \hat{y}_t^* \left[1 + \sum_{j=1}^{\infty} k_{2j}(t) a_j(t) \right]. \tag{10}$$

The second noncentral moment is given by

$$\begin{aligned} E(y_t^{*2} | \mathcal{F}_T) &= \hat{y}_t^{*2} E \left[1 + \sum_{j=1}^{\infty} k_j(t) (u_t^* - \tilde{u}_t^*)^j \right]^2 \\ &= \hat{y}_t^{*2} E \left[1 + \sum_{j=1}^{\infty} k_j^2(t) (u_t^* - \tilde{u}_t^*)^{2j} + 2 \sum_{j=1}^{\infty} \sum_{r>j} k_j(t) k_r(t) (u_t^* - \tilde{u}_t^*)^{j+r} \right. \\ &\quad \left. + 2 \sum_{j=1}^{\infty} k_j(t) (u_t^* - \tilde{u}_t^*)^j \right] \tag{11} \end{aligned}$$

Taking the expectation of both sides with respect to the Gaussian density $f(u_i^*|\mathcal{F}_T)$ we obtain:

$$E(y_i^*|\mathcal{F}_T) = \hat{y}_i^{*2} \left[1 + \sum_{j=1}^{\infty} k_j^2(t) a_j(t) + 2 \sum_{j=1}^{\infty} \sum_{r=1}^{\infty} k_j(t) k_{j+2r}(t) a_{j+r}(t) + 2 \sum_{j=1}^{\infty} k_{2j}(t) a_j(t) \right].$$

On the other hand, the square of the first moment can be written as:

$$[E(y_i^*|\mathcal{F}_T)]^2 = \hat{y}_i^{*2} \left[1 + \left(\sum_{j=1}^{\infty} k_{2j}(t) a_j(t) \right)^2 + 2 \sum_{j=1}^{\infty} k_{2j}(t) a_j(t) \right]$$

After some manipulations we obtain that:

$$\text{var}(y_i^*|\mathcal{F}_T) = \hat{y}_i^{*2} \left[\sum_{j=1}^{\infty} k_j^2(t) a_j(t) + 2 \sum_{j=1}^{\infty} \sum_{r=1}^{\infty} k_j(t) k_{j+2r}(t) a_{j+r}(t) - \left(\sum_{j=1}^{\infty} k_{2j}(t) a_j(t) \right)^2 \right] \quad (12)$$

ACKNOWLEDGEMENTS

The authors are grateful to A.C. Atkinson, W.R. Bell, V.M. Guerrero, and two anonymous referees for critically reading the manuscript and providing valuable suggestions which have led to an improvement in both the presentation and the content of the article.

NOTES

1. We adopt the convention that when $j = 1$ the product in brackets in k_j equals 1, $\prod_{i=1}^0 x_i = 1$.

Corresponding authors: Tommaso Proietti, Dipartimento S.E.F. e ME.Q., Via Columbia 2, 00133 Rome, Italy. E-mail: tommaso.proietti@uniroma2.it; Marco Riani, Dipartimento di Economia, Via Kennedy, 6 - 43100 Parma, Italy. E-mail: mriani@unipr.it

REFERENCES

ATKINSON, A. C. (1985) *Plots, Transformations and Regression: An Introduction to Graphical Methods of Diagnostic Regression Analysis*. New York: Oxford University Press.
 ATKINSON, A. C. and RIANI, M. (2000) *Robust Diagnostic Regression Analysis*. New York: Springer Verlag.

- ATKINSON, A. C. and RIANI, M. (2002) Tests in the fan plot for robust, diagnostic transformations in regression, *Chemometrics and Intelligent Laboratory Systems* 60, 87–100.
- BELL, W. R. and HILLMER, S. C. (1984) Issues involved with the seasonal adjustment of economic time series. *Journal of Business and Economic Statistics* 2, 291–320.
- BOX G. E. P. and COX, D. R. (1964) An analysis of transformations (with discussion). *Journal of the Royal Statistical Society B* 26, 211–46.
- BOX G. E. P. and JENKINS, G. M. (1973) Some comments on a paper by Chatfield and Prothero and on a review by Kendall. *Journal of the Royal Statistical Society A* 136, 337–52.
- CARROLL, R. J. and RUPPERT, D. (1991) Prediction and tolerance intervals with transformation and/or weighting. *Technometrics* 33, 197–210.
- CHATFIELD, C. and PROTHERO, D. L. (1973) Box-Jenkins seasonal forecasting: problems in a case-study. *Journal of the Royal Statistical Society A* 136, 295–315.
- CLEVELAND, W. S., DUNN, D. M. and TERPENNING, I. J. (1978) SABL – a resistant seasonal adjustment procedure with graphical methods for interpretation and diagnosis. In *Seasonal Analysis of Economic Time Series* (ed. A. ZELLNER). Washington, D.C.: U.S. Department of Commerce, Bureau of the Census, pp. 201–31.
- COOK, R. D. and WEISBERG, S. (1999) *Applied Regression Including Computing and Graphics*. New York: Wiley.
- DOORNIK, J. A. (2001) *Ox 3.0 – An Object-Oriented Matrix Programming Language*. London: Timberlake Consultants Ltd.
- DURBIN, J. and KOOPMAN, S. J. (2001). *Time Series Analysis by State Space Methods*. Oxford: Oxford University Press.
- DURBIN, J. and KOOPMAN, S. J. (2002) A simple and efficient simulation smoother for state space time series analysis. *Biometrika* 89, 603–15.
- FINDLEY, D. F., MONSELL, B. C., BELL, W. R., OTTO, M. C. and CHEN, B. C. (1998) New capabilities and methods of the X-12-ARIMA seasonal adjustment program (with Discussion). *Journal of Business and Economic Statistics* 16, 127–77.
- GÓMEZ, V. and MARAVALL, A. (1997) *Programs TRAMO and SEATS: Instructions for the User (Beta Version: June 1997)*, Working Paper 97001. Madrid: Ministerio de Economía y Hacienda, Dirección General de Análisis y Programación Presupuestaria.
- GRANGER, C. W. J. and NEWBOLD, P. (1976) Forecasting transformed time series. *Journal of the Royal Statistical Society B* 38, 189–203.
- GUERRERO, V. M. (1993) Time-series analysis supported by power transformations. *Journal of Forecasting* 12, 37–48.
- HARVEY, A. C. (1989) *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge: Cambridge University Press.
- HARVEY, A. C. and PROIETTI, T. (2005) *Readings in Unobserved Components Models, Advanced Texts in Econometrics*. Cambridge: Oxford University Press.
- JOHN J. A. and DRAPER N. R. (1980) An alternative family of transformations. *Applied Statistics* 29, 190–7.
- DE JONG, P. and SHEPHARD, N. G. (1995) The simulation smoother for time series models. *Biometrika* 82, 339–50.
- KOOPMAN, S. J., SHEPHARD, N. and DOORNIK, J. A. (1999) Statistical algorithms for models in state space using SsfPack 2.2. *Econometrics Journal* 2, 113–66.
- MARAVALL, A. (1985) On structural time series models and the characterization of components. *Journal of Business & Economic Statistics* 3, 350–55.
- NEYMAN, J. N. and SCOTT, E. L. (1960) Correction for bias introduced by a transformation of variables. *The Annals of Mathematical Statistics* 31, 643–55.
- PANKRATZ, A. and DUDLEY, U. (1987) Forecasts of power-transformed series. *Journal of Forecasting* 6, 239–48.
- PIESSENS, R., DE DONCKER-KAPENGA, E., UBERHUBER, C. and KAHANER, D. (1983) *QUADPACK: A Subroutine Package for Automatic Integration*. Berlin: Springer.
- PRESS, W. H., FLANNERY, B. P., TEUKOLSKY, S. A. and VETTERLING, W. T. (1992) *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge: Cambridge University Press.
- PROIETTI, T. (2000) Comparing seasonal components for structural time series models. *International Journal of Forecasting* 16, 247–60.
- PROIETTI, T. and RIANI, M. (2007) Transformations and seasonal adjustment: analytic solutions and case studies. London: Department of Statistics, School of Economics. Research report no. 143.

- RIANI, M. (1998) Weights and robustness of model based seasonal adjustment. *Journal of Forecasting* 17, 19–34.
- RIANI, M. (2004) Extensions of the forward search to time series. *Studies in Non Linear Dynamics and Econometrics* 8, 1–23.
- TAYLOR, J. M. G. (1986) The retransformed mean after a fitted power transformation. *Journal of the American Statistical Association* 81, 114–8.
- THOMSON, P. and OZAKI, T. (2002) Transformation and trend-seasonal decomposition, Invited paper. *Proc. 3rd Int. Symp. Frontiers of Time Series Modeling*. Tokyo, Japan, 197–212.
- TUKEY, J. W. (1957) On the comparative anatomy of transformations. *Annals of Mathematical Statistics* 28, 602–32.
- YEO, I. K. and JOHNSON, R. A. (2000) A new family of power transformation to improve normality or symmetry, *Biometrika* 87, 954–9.