Time series-dependent selection of an appropriate seasonal adjustment approach

Karsten Webel*
Deutsche Bundesbank, Frankfurt am Main, Germany
E-mail: karsten.webel@bundesbank.de

Abstract

When choosing a software package for conducting seasonal adjustment as part of their daily routines, decisions of many statistical agencies are based on pragmatic reasons, such as employees' individual backgrounds, data users' demands (e.g. for low revisions) and the program's suitability for statistical mass production. Then, usually all time series, or at least broad subsets thereof, are seasonally adjusted according to the approach implemented in the software package chosen. Recent releases of X-13ARIMA-SEATS and Demetra+ may change habits as these programs include both the nonparametric X-11 approach and the parametric ARIMA model-based approach. Hence, users may choose between both methods for each particular time series under review. Accordingly, the question immediately arises which criteria one should rely on when making this selection. We develop a decision tree that combines theoretical considerations regarding conceptual differences between both approaches with empirical findings. In particular, the latter include visual inspection of squared gains of X-11 and SEATS seasonal adjustment filters as well as calculation of diverse revision measures, which we demonstrate for selected German key macroeconomic indicators.

Keywords: Demetra+, signal extraction, unobserved components, X-13ARIMA-SEATS.

1. Introduction

Both X-13ARIMA-SEATS (X13AS) and Demetra+ merge capabilities of a nonparametric and a parametric approach to seasonal adjustment (SA) since they incorporate the X-11 and SEATS core developed originally by Shiskin et al. (1967) and Gómez and Maravall (1996), respectively. Users may thus choose between these methods without switching between different software packages. This immediately raises the question of which criteria one should base this choice on. For that purpose, we suggest a decision tree that relies on both theoretical and empirical issues.

The remainder of this paper is organized as follows. Section 2 briefly outlines basic concepts of SA. Section 3 focuses on those branches of our decision tree that exploit conceptual differences between the X-11 and ARIMA model-based (AMB) method. Complementary, Section 4 illustrates its rather advanced empirical branches by comparing performances of both algorithms using monthly industrial output and orders received by industry. Finally, section 5 concludes.

2. Basic theory

Let $\{x_t\}$ denote the time series under review and assume that it admits a signal-plus-noise representation according to

$$x_t = s_t + n_t, \tag{1}$$

where the signal $\{s_t\}$ and the noise $\{n_t\}$ are orthogonal as they shall capture all non-seasonal and seasonal variation, respectively, of $\{x_t\}$. Hence, SA is basically a signal extraction problem which can be solved by linear filtering. To see this, let $\boldsymbol{x} =$

---

* This paper represents the author's personal opinion and does not necessarily reflect the views of the Deutsche Bundesbank or its staff. The author thanks Robert Kirchner and Jens Mehrhoff for valuable comments. All remaining errors are, of course, the author's sole responsibility.

$(x_1, \dots, x_T)'$ denote a finite realization of $\{x_t\}$. We may then express any estimator of the SA time series as

$$\hat{\boldsymbol{s}} = \boldsymbol{W} \cdot \boldsymbol{x}, \tag{2}$$

where $\hat{\boldsymbol{s}} = (\hat{s}_1, \dots, \hat{s}_T)'$ and $\boldsymbol{W}$ is a matrix of order $T \times T$ whose t-th row, $\boldsymbol{w}_t$, contains the filter weights employed to estimate s_t. For each $t \in \{1, \dots, T\}$, we may thus rewrite (2) as

$$\hat{s}_t = \boldsymbol{w}_t \cdot \boldsymbol{x} = \sum_{j=-(T-t)}^{t-1} w_{t,j} \cdot x_{t-j}. \tag{3}$$

It becomes obvious from (3) that the filter weights stored in $\boldsymbol{W}$ do not only depend on t but also on the number of observations available, T. Assuming for simplicity that the latter is odd, i.e. $T = 2\tau + 1$ for some $\tau \in \mathbb{N}$, the symmetric central SA filter $\boldsymbol{w}_{\tau+1}$ and the asymmetric concurrent SA filter $\boldsymbol{w}_T$ are usually of greatest interest.

Both SA methods implemented in X13AS differ substantially in the way these weights are derived. In essence, the X-11 approach relies on an iterated application of various predefined trend and (period-specific) seasonal filters. Both types of moving averages can be selected automatically according to some built-in heuristic criteria or manually to match the user's needs. Either way, final X-11 filters result from the convolution of all trend and seasonal filters chosen during each iteration step. Their weights can be easily determined from the impulse response method described in Section 3.4 of Ladiray and Quenneville (2001). In practice, however, $\boldsymbol{x}$ is usually extended with ARIMA forecasts to mitigate boundary problems at its current end. In this case, Equation (2) and (3), respectively, can be modified straightforwardly to include these forecasts.

In contrast, the AMB approach assumes that the dynamics of both unobserved components of model (1) can be adequately captured by individual ARIMA processes and, thus, $\boldsymbol{x}$ admits an ARIMA representation by construction. In practice, the latter is estimated first and individual ARIMA models of the components are derived from it afterwards according to a complex decomposition algorithm which might also include certain approximations in case no admissible decomposition can be found. Eventually, this procedure yields the data-driven Wiener-Kolmogorov (WK) filter for the signal of model (1) which automatically defines the weights of SEATS's SA filters. Hence, the SA time series' estimator found by the AMB method is MSE-optimal by construction.

Further details on both approaches are provided by Findley et al. (1998), Gómez and Maravall (2001), Ladiray and Quenneville (2001) and Maravall (1995), inter alia.

3. First empirical considerations

Figure 1 presents our first ideas on a decision tree for choosing between the X-11 and AMB algorithm. As shown in its top dashed rectangle, a selection may be made in some situations by exploiting fundamental differences between the philosophies underlying both methods. In particular, SA with X-11 has advantages whenever AMB theory's basic assumption of the data generating process coinciding with an ARIMA process is likely to be too restrictive to capture all dynamics of the time series observed. This, for example, might be the case if the latter exhibits seasonal heteroskedasticity as ARIMA models cannot reproduce any GARCH-type behaviour. In contrast, such effects can be taken into account, at least up to some extent, within X-11's built-in extreme value detection procedure by considering period-specific variances of the signal's irregular component, provided their importance is sufficiently high. If not, usage of the X-11 or AMB approach does not really make a difference as far as seasonal heteroskedasticity is concerned, at least according to our experience.

Another situation where reliance on the latter SA method might be less preferable emerges when major causes of (moving) seasonality in observed unadjusted figures are substantially different across periods. For example, in German retail sales the usual peak in December attributable to the Christmas business has decreased continuously over years. One major reason is that Christmas bonuses have been paid by fewer firms

The choice between the X-11 and AMB approach to SA — Figure 1

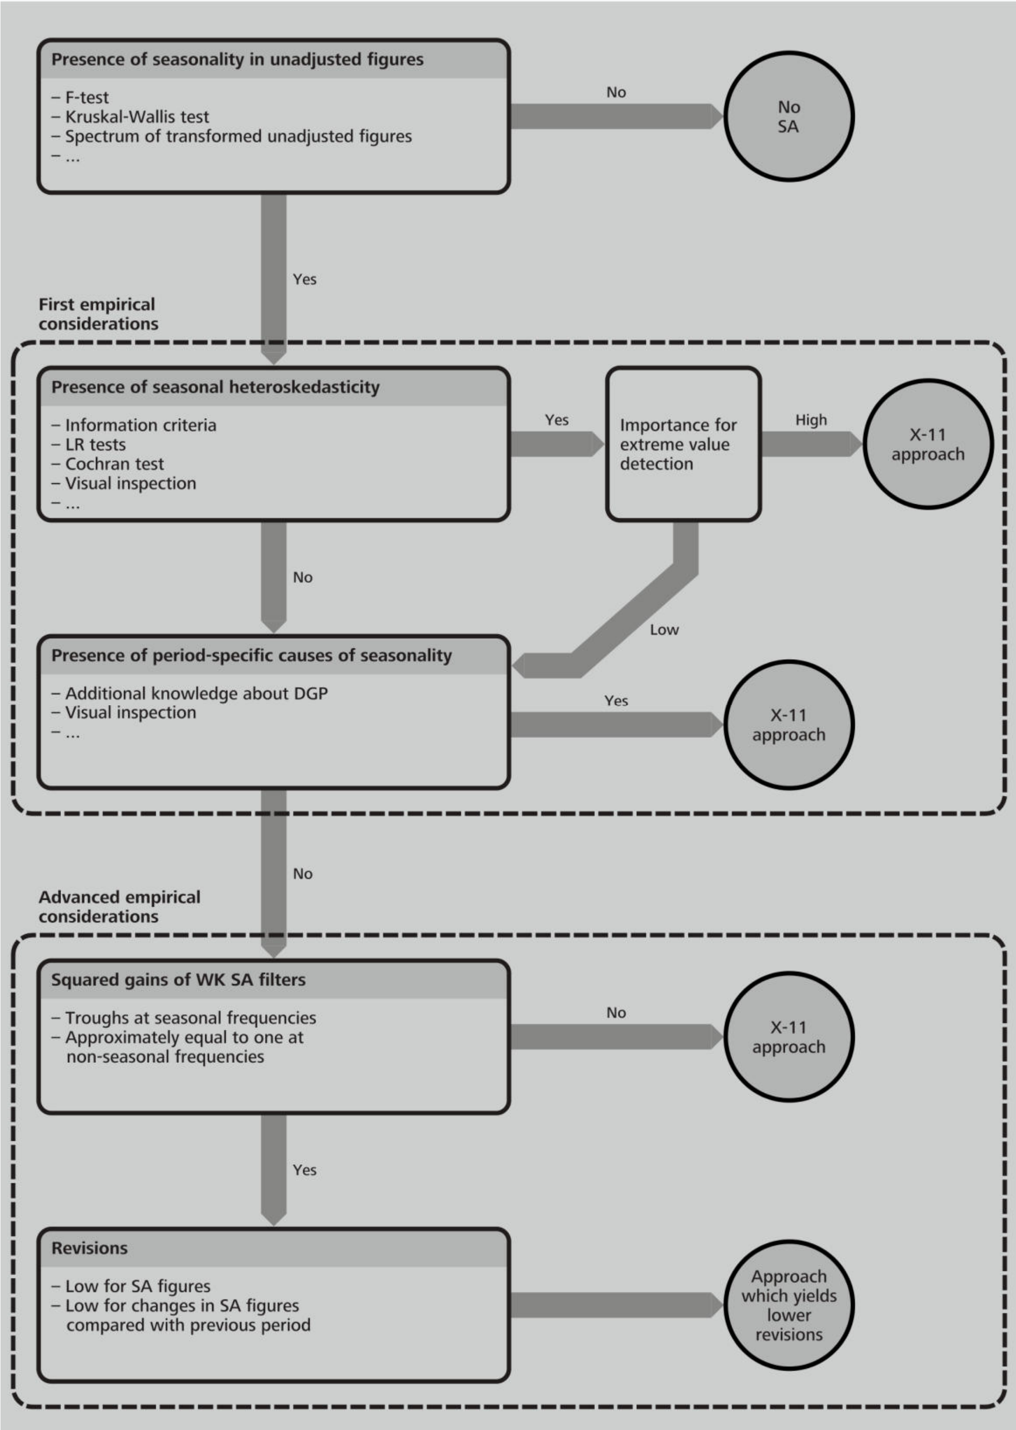

Deutsche Bundesbank

or have been cut steadily compared to regular salaries. Unadjusted figures for December should thus be treated somewhat individually by, for instance, application of shorter seasonal filters in X-11 as compared to other months that should be barely affected by this change. In contrast, such a capability is not part of the AMB framework. Of course, these first empirical considerations do usually not suffice to find an appropriate SA method for time series that behave rather less exceptionally. Instead, as shown in the bottom dashed rectangle of Figure 1, our decision tree needs to be extended with some advanced empirical branches in these cases.

RegARIMA modelling Table 1
Industry, 2005 = 100, calendar adjusted

Time series	ARIMA model	Parameter estimates [1)] Non-seasonal	Seasonal
Output	(3 1 1)(0 1 1)	0.3150 (0.1957) 0.1317 (0.0835) 0.2102 (0.0618) 0.6528 (0.1952)	0.5907 (0.0514)
Orders received	(1 1 0)(0 1 1)	-0.2958 (0.0589)	0.8013 (0.0410)

1 In the order of appearance in the ARIMA model, standard errors in parentheses.

4. Advanced empirical considerations

We consider monthly output in industry at constant prices and monthly orders received by industry at current prices as of January 1991 up to December 2012 with respective annual averages of 2005 being set to 100. Both time series are well-behaved in the sense that they pass our first empirical considerations above. Also, they have already been adjusted for calendar effects using linear regression techniques. We run X13AS with default options, using its Windows companion Win X-13 (version 1.0 build 150). In addition, R (version 2.14.1) is employed to calculate squared gains of X-11 filters, as these are not provided by any Win X-13 output.

In the first step, we fit a regARIMA model to the log of each indicator to account for outliers and to obtain two years of forecasts. Thereby, the log transformation is preferred by the automatic log-level test of X13AS. Similarly, final ARIMA model choices are based on the program's automatic model selection procedure. Results are summarized in Table 1, where the standard notation $(p\ d\ q)(P\ D\ Q)$ is used to indicate an ARIMA model whose orders of the non-seasonal AR, differencing and MA operator and given by p, d and q, respectively. Likewise, capitol letters denote the orders of corresponding seasonal operators.

In a second step, we compare squared gains of SA filters chosen in X-11 and SEATS. The basic idea is that any "acceptable" SA filter should completely eliminate seasonal variation without altering movements associated with non-seasonal frequencies. Assuming stable seasonality, the squared gain of an ideal SA filter thus equals one at all frequencies except the seasonal ones where it admits discontinuities as it drops to zero. In practice, of course, squared gains of SA filters are certain to deviate from this ideal due to the finite length of any time series observed and, consequently, their extent of deviation is crucial to selecting an appropriate SA approach. For the sake of brevity, we restrict ourselves to central SA filters. Respective results, e.g. for concurrent SA filters, are available from the author upon request. Regarding X-11, we further distinguish between an automatic run of X13AS (X-11-AFS) where trend and seasonal filters are chosen according to the automatic selection procedures implemented and an official run (X-11-OFC) which closely mimics the specifications used within the production process of official seasonally adjusted figures.[1] In particular, the latter makes use of period-specific seasonal filters which is not possible when relying on an automatic filter selection. Eventually, squared gains of central WK filters are directly provided by X13AS.

The top panel of Figure 2 shows the squared gains of the X-11 and SEATS filters considered for SA of industrial output. Both X-11 filters perform reasonable well as their squared gains stay close to one before the first seasonal frequency and decrease rapidly to zero in the vicinity of seasonal frequencies. Thereby, dips are slightly narrower

[1] Note, however, that macroeconomic aggregates as well as their major components are usually SA according to an indirect approach in official statistics. In addition, SA figures obtained from the direct approach are typically considered for comparative analyses.

Squared gain of central SA filter Figure 2

Industry, 2005 = 100, calendar adjusted

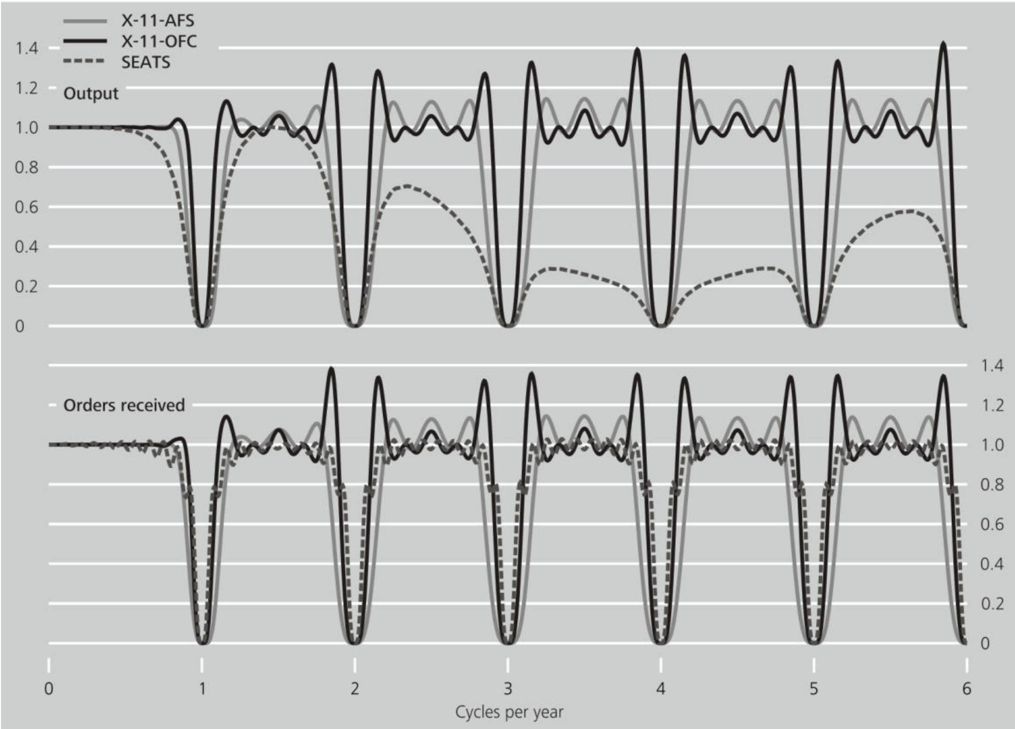

Deutsche Bundesbank

for the official run. In addition, both squared gains display oscillatory behavior between seasonal frequencies typical of finite filters. As a consequence, variation in unadjusted output associated with near-seasonal frequencies might be amplified in SA output by a factor slightly above 1.4 at maximum. In contrast, the WK filter performs rather poorly, even though it still removes virtually all seasonality contained in unadjusted output. Since its squared gain, however, does not manage to exceed 0.7 behind the second seasonal frequency and, moreover, stays below 0.3 between the third and fifth one, the WK filter also eliminates much variation not associated with seasonal behavior, resulting in SA output that is much smoother than the one obtained from the X-11 approach. Since SA should eliminate only seasonal behaviour by definition, the latter method appears preferable for industrial output if judgment is based on squared gains of SA filters. Thereby, SA figures do not significantly depend on whether the automatic or official run of X-11 is employed.

When looking at orders received, things are completely different. As shown in the bottom panel of Figure 2, SA with X-11 once again provides acceptable results. Squared gains of both SA filters considered display behaviour very similar to that observed for output and, thus, respective comments made above apply basically in the same way. However, X-11 is outperformed by SEATS because its squared gain's oscillations between seasonal frequencies are significantly lower and, thus, non-seasonal intra-year variation is less amplified in SA orders received if SEATS is applied instead of X-11. In sum, the AMB method seems to be advantageous for orders received so far. Nevertheless, we shall also compare both SA approaches with respect to revisions before making a final decision since this advantage is rather slight than huge.

In a third step, we therefore calculate diverse revision measures for SA orders received. Since we are primarily interested in the way revisions are induced by application of a particular SA method, we do not consider them in real-time, excluding corrections of unadjusted figures from our analysis. Instead, we seasonally adjust truncated versions of orders received to generate revisions in its SA figures and their

Revision analysis of orders received Table 2
January 2000 – December 2007, data as of February 2013

SA core	SA time series [1] MR	MAR	STD	Changes in SA time series compared with previous period [2] MR	MAR	STD
X-11-AFS	0.055	0.494	0.633	0.035	0.582	0.736
X-11-OFC	0.080	0.466	0.594	0.043	0.593	0.736
SEATS	0.017	0.399	0.543	0.014	0.462	0.581

1 As a percentage. **2** In percentage points.

month-on-month changes. Thereby, the ARIMA model is re-estimated for each truncated series. Table 2 reports mean revisions (MR), mean absolute revisions (MAR) and the revisions' standard deviations (STD) that result from both SA approaches. These measures tend to be slightly smaller for the AMB method, while being basically of the same size for both X-11 runs. Combining information obtained from squared gains and revisions, SA with SEATS seems thus to be appropriate for orders received. However, it should be noted eventually that for this indicator both methods produce SA figures that are not substantially different from each other. For example, the mean difference and mean absolute difference between SA orders received obtained from the official X-11 run and SEATS amount to 0.01 and 0.24 index point, respectively.

5. Summary

According to our decision tree, the X-11 approach is preferred for industrial output, whereas SA with SEATS is recommended for orders received by industry. These choices might be explained by the fact that AMB theory works extremely well for ARIMA models that are somewhat close to Airline models. In contrast, the X-11 method might be advantageous for time series fitted by ARIMA models of rather complex structure, where the AMB algorithm may run into problems of non-admissibility. However, this is rather a conjecture than an explanation and demands further research. For example, the awkward curvature of industrial output's squared gain of the central SEATS SA filter may result from transitory effects falsely attributed to the seasonal component. To verify this, SA with the AMB method should be rerun with user-specific options instead of default ones. In addition, future empirical research should also put more emphasis on squared gains of concurrent SA filters and incorporate phase delays into the decision-making process.

References

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," *Journal of Business & Economic Statistics*, 16, 127-152.

Gómez, V. and Maravall, A. (1996) "Programs TRAMO and SEATS – Instructions for the User," *Working Paper No 9628*, Banco de España, Madrid.

Gómez, V. and Maravall, A. (2001) "Seasonal Adjustment and Signal Extraction in Economic Time Series," In: Peña, D., Tiao, G. C. and Tsay, R. S. (Eds.) *A Course in Time Series Analysis*, 202-247, Wiley, New York.

Ladiray, D. and Quenneville, B. (2001) *Seasonal Adjustment with the X-11 Method*, Lecture Notes in Statistics, 158, Springer, New York.

Maravall, A. (1995) "Unobserved Components in Economic Time Series," In: Pesaran, M. H. and Wickens, M. R. (Eds.) *The Handbook of Applied Econometrics*, 1, 12-72, Blackwell Publishing, Oxford.

Shiskin, J., Young, A. H. and Musgrave, J. C. (1967) "The X-11 Variant of the Census Method II Seasonal Adjustment Program," *Technical Paper No 15*, U. S. Department of Commerce, Bureau of the Census, Washington, D. C.