Discussion Paper

Deutsche Bundesbank
No 41/2018

Seasonal adjustment of daily time series

Daniel Ollech

Deutsche Bundesbank, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main,
Postfach 10 06 02, 60006 Frankfurt am Main

Tel +49 69 9566-0

Please address all orders in writing to: Deutsche Bundesbank,
Press and Public Relations Division, at the above address or via fax +49 69 9566-3077

Internet http://www.bundesbank.de

ISBN 978–3–95729–510–1 (Printversion)
ISBN 978–3–95729–511–8 (Internetversion)

Non-technical summary

Research Question

Currently, the methods used by producers of official statistics do not facilitate the seasonal and calendar adjustment of time series with daily observations. However, an increasing number of possibly seasonal series with a weekly or daily periodicity becomes available. Here, we develop a procedure to estimate and adjust for periodically recurring systematic effects and the influence of moving holidays in time series with daily observations.

Contribution

Daily time series can contain three different seasonalities. Exemplarily, the currency in circulation is influenced by a weekday effect as the demand for banknotes increases towards the weekend. Monthly recurring troughs and peaks can often be observed, because salary payments tend to be concentrated around the turn of the month. Due to an increase in consumption the demand for currency reaches its high around Christmas time. In addition to these seasonal influences, the currency in circulation is positively impacted by Easter and other public holidays whose date changes from year to year. The procedure for daily calendar and seasonal adjustment (DSA) introduced in this paper combines a local regression based iterative seasonal adjustment routine with a time series regression model for the estimation of moving holiday and outlier effects.

Results

To assess the validity of the developed approach, the DSA adjusted time series are transformed to a monthly basis and then compared to results obtained by established procedures for monthly data for the German currency in circulation as well as a for a set of simulated time series. The results show that all procedures estimate similar seasonally adjusted series. The comparison with the established seasonal adjustment routines indicates a high validity of the chosen approach. Thus, the DSA procedure closes a gap by facilitating the seasonal and calendar adjustment of daily time series.

Nichttechnische Zusammenfassung

Fragestellung

Die in der offiziellen Statistik verwendeten Methoden erlauben derzeit keine Saisonbereinigung von Zeitreihen mit täglichen Beobachtungen. Gleichwohl steigt die Anzahl verfügbarer höherfrequenter Zeitreihen mit saisonalem Verlauf. Das Ziel der vorliegenden Untersuchung ist die Entwicklung eines Verfahrens zur Schätzung und Bereinigung periodisch wiederkehrender, systematischer Effekte und des Einflusses beweglicher Feriertage in Zeitreihen mit täglichen Beobachtungen.

Beitrag

Tägliche Zeitreihen können drei unterschiedliche, sich überlagernde Saisonmuster enthalten. Beispielsweise steigt der Banknotenumlauf aufgrund vermehrter Einkäufe üblicherweise jeweils besonders stark zum Wochenende. Ein monatlich wiederkehrendes Muster ist insbesondere aufgrund von Lohn- und Gehaltszahlungen um den Monatswechsel herum identifizierbar. Im Jahresverlauf sind die deutlichsten Spitzen in der Banknotennachfrage infolge eines gesteigerten Konsumverhaltens rund um Weihnachten zu beobachten. Zusätzlich zu diesen saisonalen Regelmäßigkeiten wird der Banknotenumlauf durch die Lage beweglicher Feiertage wie Ostern beeinflusst. Das hier vorgestellte Verfahren zur Kalender- und Saisonbereinigung täglicher Daten (DSA) kombiniert eine mehrstufige, auf lokalen Regressionen basierende Saisonbereinigung mit an die Eigenschaften von Zeitreihen angepassten Regressionsmodellen zur Schätzung kalendarischer und Ausreißereffekte.

Ergebnisse

Zur Überprüfung der Validität des Ansatzes werden die ermittelten kalender- und saisonbereinigten Zeitreihen in monatliche Zeitreihen transformiert und mit den Ergebnissen etablierter Verfahren zur monatlichen Saisonbereinigung verglichen. Als Grundlage dient hierbei neben dem deutschen Banknotenumlauf ein Datensatz simulierter Zeitreihen. Die Gegenüberstellung zeigt, dass alle Verfahren vergleichbare saisonbereinigte Reihen schätzen. Auch berechnete Fehlermaße weisen auf eine qualititative Ähnlichkeit hin. Die Vergleiche mit den Ergebnissen der etablierten Saisonbereinigungsmethoden deuten auf eine hohe Validität des gewählten Ansatzes. Das DSA-Verfahren lässt sich demnach zur Saisonbereinigung täglicher Daten verwenden.

Deutsche Bundesbank Discussion Paper No 41/2018

Seasonal Adjustment of Daily Time Series*

by Daniel Ollech

Abstract

Currently, the methods used by producers of official statistics do not facilitate the seasonal and calendar adjustment of daily time series, even though an increasing number of series with daily observations are available. The aim of this paper is the development of a procedure to estimate and adjust for periodically recurring systematic effects and the influence of moving holidays in time series with daily observations. To this end, an iterative STL based seasonal adjustment routine is combined with a RegARIMA model for the estimation of calendar and outlier effects. The procedure is illustrated and validated using the currency in circulation in Germany and a set of simulated time series. A comparison with established methods used for the adjustment of monthly data shows that the procedures estimate similar seasonally adjusted series. Thus, the developed procedure closes a gap by facilitating the seasonal and calendar adjustment of daily time series.

Keywords: Seasonal adjustment; STL; Daily time series; Seasonality.

JEL classification: C14, C22, C53.

*Contact address: Daniel Ollech, Deutsche Bundesbank, Central Office, Statistics Department, Wilhelm-Epstein-Strasse 14, 60431 Frankfurt am Main, Germany. Phone: +49 69 9566 6250. E-mail: daniel.ollech@bundesbank.de. The author thanks Julian LeCrone, Karsten Webel, Robert Kirchner, Edgar Brandt, Andreas Lorenz, Andreas Dietrich, Susanne Stollenmayer, Duncan Elliott, Christoph Hanck, and conference participants at the ISF 2016 in Santander and CFE 2016 in Seville for helpful comments. Discussion Papers represent the authors' personal opinions and do not necessarily reflect the views of the Deutsche Bundesbank or the Eurosystem.

1 Introduction

Progress in information technology will increase the availability of time series with a daily or even higher periodicity that exhibit seasonal behaviour. Possible examples range from data on air pollution, web-search keywords, and road traffic data to economic variables such as online sales, retail prices, and the amount of banknotes in circulation. Often, new developments cannot be evaluated directly from these unadjusted data, as periodical influences may obscure underlying tendencies. Furthermore, periodic influences may be of interest themselves. An increasing number of banknotes towards the end of the year may be due to the annually recurring Christmas shopping related demand spikes instead of worrisome money hoarding. Seasonal adjustment aims at detecting and eliminating these kinds of periodic influences to enhance the interpretability of the data.

Currently, there exists no officially recommended method for seasonal and calendar adjustment of time series with daily observations (cf Eurostat, 2015). The aim of this paper is to identify a suitable procedure to estimate and adjust for regularly and periodically recurring systematic effects as well as the influence of moving holidays for daily time series. The procedure should be flexible enough to estimate all seasonal patterns with different periodicities. At the same time, it should be sufficiently fast to allow the adjustment of a large number of series on a daily basis.

In their implementation in various software packages the well-known seasonal adjustment procedures TRAMO/SEATS and Census X-13ARIMA-SEATS (X-13)[1] do not allow for the adjustment of data with a higher than monthly frequency. Alternative methods have been developed that aim at estimating some or all of the seasonal effects in higher frequency time series. One of the first method's proposed was Akaike's BAYSEA procedure that is based upon Bayesian time series modelling and was developed as an alternative to X-11 (Akaike, 1980; Young, 1996). In theory, the approach is not restricted to monthly or quarterly observations. Yet, it can be demonstrated by simulation that BAYSEA does not yield plausible results for higher frequency data with a substantial intra-annual seasonal pattern, due to very slow convergence (see appendix A).

Pierce, Grupe, and Cleveland (1984) use a two-step model-based approach to seasonally adjust weekly monetary aggregates. Trigonometric functions are designed to capture any deterministic weekly seasonality and an ARIMA model is fitted to the residuals to obtain estimates for the remaining stochastic seasonality. Without further adaptations, simple seasonal ARIMA models do not capture all different seasonal frequencies that we observe for daily time series. Especially the simultaneous modelling of weekly and annually recurring patterns is challenging, as will be discussed later.

[1] These are implemented in X-13ARIMA-SEATS, TSW rev. 924 and JDemetra+ 2.2.

Dokumentov and Hyndman (2015) develop a regression based seasonal adjustment procedure, STR, that aims at providing automatic detection of multiple relevant seasonal cycles. However, the user has only limited influence on the outcome of this procedure. Depending on the complexity of the time series, the adjustment may also be computationally expensive. Section 4.2 includes an empirical comparison between STR and the routine developed in this paper.

The software STAMP 8.2, which is designed to seasonally adjust data using structural time series modelling, assumes that daily time series have an intra-weekly pattern only (see Koopman, Harvey, Doornik, and Shephard, 2009). Harvey, Koopman, and Riani (1997) employ structural models to adjust weekly monetary aggregates. For intra-monthly and intra-annual effects they suggest using either trigonometric or spline functions. Koopman and Ooms (2003, 2006) show that state space modelling can be used to analyse seasonal daily time series as well, though their main focus lies on short-term forecasting instead of seasonal adjustment as such.

Semi-parametric STL provides another method, which is based on a locally weighted regression smoother (Loess). Compared to the methods discussed above, its main advantage is its high flexibility with respect to the frequency of the time series (Cleveland, Cleveland, McRae, and Terpenning, 1990). Additionally, the fast computation of the STL algorithm makes it feasible to adjust different seasonal frequencies in an integrated iterative framework. Other authors have used STL for different applications involving seasonal time series. Verbesselt, Hyndman, Newnham, and Culvenor (2010) develop a trend and seasonal change detection method that is based on STL and can be used for daily and other higher frequency time series. Thus, STL has properties that make it a suitable candidate for the seasonal adjustment of data with higher frequency.

We contribute to the existing literature by devising a seasonal adjustment routine for daily data. In addition to STL's limitations to single seasonal frequencies it cannot deal with calendar effects such as the influences of moving holidays. Therefore, we use STL as the basis of the seasonal adjustment and combine it with RegARIMA based estimations of calendar effects to obtain the daily seasonal adjustment (DSA) procedure.

The remainder of the paper is structured as follows: Section 2 presents key aspects of STL. Section 3 adapts STL to handle daily time series. Section 4 applies the procedure to the daily currency in circulation in Germany and to a set of simulated time series and then compares the results to those obtained from established seasonal adjustment procedures for monthly data. Section 5 concludes.

2 STL

Based on the well-known component model

$$Y_t = T_t + S_t + I_t \quad \forall \quad t = 1, ..., N \tag{1}$$

STL decomposes a time series (Y_t) into a trend-cycle (T_t), a seasonal (S_t) and an irregular component (I_t) using Loess regressions and moving averages (Cleveland et al., 1990). In Loess regressions, a weight is attached to each observation of the time series (Cleveland and Devlin, 1988). This weight is negatively related to the distance (in time) between a given observation and the value that is to be smoothed. If the distance is too large, the weight is zero. Thus, Loess regressions are local regressions because each value is regressed on a local neighbourhood of a linear or quadratic function of the (weighted) observations. For any point in time t, the weight of the observation at t_i is given by

$$v_i(t) = \left[1 - \left(\frac{|t_i - t|}{\delta_\gamma(t)}\right)^3\right]^3 \tag{2}$$

with $\delta_\gamma(t) = |t_i - t_\gamma|$ being the distance between the γ^{th} farthest t_i and t. The smoothing parameter γ determines the number of neighbouring observations included in the local regression. In STL, the user can specify the value of γ for the computation of the trend and seasonal component. The flexibility of the identified seasonal factor decreases with a higher value for γ. Correspondingly, the value of γ should be adapted to the variability of the observed seasonality (see also Section 3 for examples).

Essentially, STL consists of an inner loop that is used for the decomposition of the series into the trend, seasonal and irregular component, and an outer loop for extreme value adjustment (see Figure 1). The latter increases the robustness of the decomposition against outliers and thereby mitigates the influence of extraordinary events on the identification of the seasonal factors. STL can only adjust the seasonal pattern of a single frequency f at a time. The other seasonal patterns will be included in the irregular component.

2.1 Inner Loop

For the decomposition of Y_t, the inner loop goes through the following computations in the k-th iteration:

1. **Trend adjustment:**

$$Y_t - T_t^{\{k-1\}} = S_t^{\{k\}} + I_t^{\{k\}} \equiv TA_t^{\{k\}}, \text{ with } T_t^{\{0\}} = \mathbf{0}. \tag{3}$$

2. **Preliminary periodwise smoothing:** Each subseries of $TA_t^{\{k\}}$, i.e. each weekday, day of the month or day of the year for a daily time series or each month for a monthly time series, is smoothed by Loess to yield a preliminary seasonal factor $S_{pre,t}^{\{k\}}$. The γ for these regressions is the only parameter without a default in the STL procedure and thus has to be specified by the user (see Cleveland et al., 1990).

3. **Smoothing preliminary seasonal component:** A low-pass filter, composed of moving average- and Loess-filters, is used to capture any low-frequency movements $L_t^{\{k\}}$ from $S_{pre,t}^{\{k\}}$.

4. **Obtaining seasonal component:**

$$S_t^{\{k\}} = S_{pre,t}^{\{k\}} - L_t^{\{k\}}. \tag{4}$$

5. **Seasonally adjusting the original time series:**

$$Y_t^{\{k\}} - S_t^{\{k\}} \equiv SA_t^{\{k\}}. \tag{5}$$

6. **Obtaining trend:** A Loess filter is applied to $SA_t^{\{k\}}$ to yield $T_t^{\{k\}}$.

2.2 Outer Loop

Usually, only a few iterations of the inner loop are needed to ensure convergence of the results. By default, the number of iterations is set to 3 and the outer loop is omitted. Only if deemed necessary by the user, robustness weights for the Loess regressions of the inner loop are determined for each time point as a function of the size of the irregular component derived when the inner loop has reached its maximum number of iterations. The irregular component can be obtained by rearranging Equation 1:

$$I_t = Y_t - T_t - S_t. \tag{6}$$

Then, for the Loess regressions in the 2. and 6. part of the inner loop, the weights defined in Equation 2 are multiplied with the robustness weights given by Tukey's (1960) biweight function

$$\omega_t = \begin{cases} \left(1 - \left[\frac{|I_t|}{6 \cdot median(|\{I_t\}_{t=1}^N|)}\right]^2\right)^2 & \text{if } |I_t| < 6 \cdot median(|\{I_t\}_{t=1}^N|) \\ 0 & \text{else} \end{cases} \tag{7}$$

that strongly downweighs highly irregular observations.

STL Algorithm for a Given Frequency

Figure 1

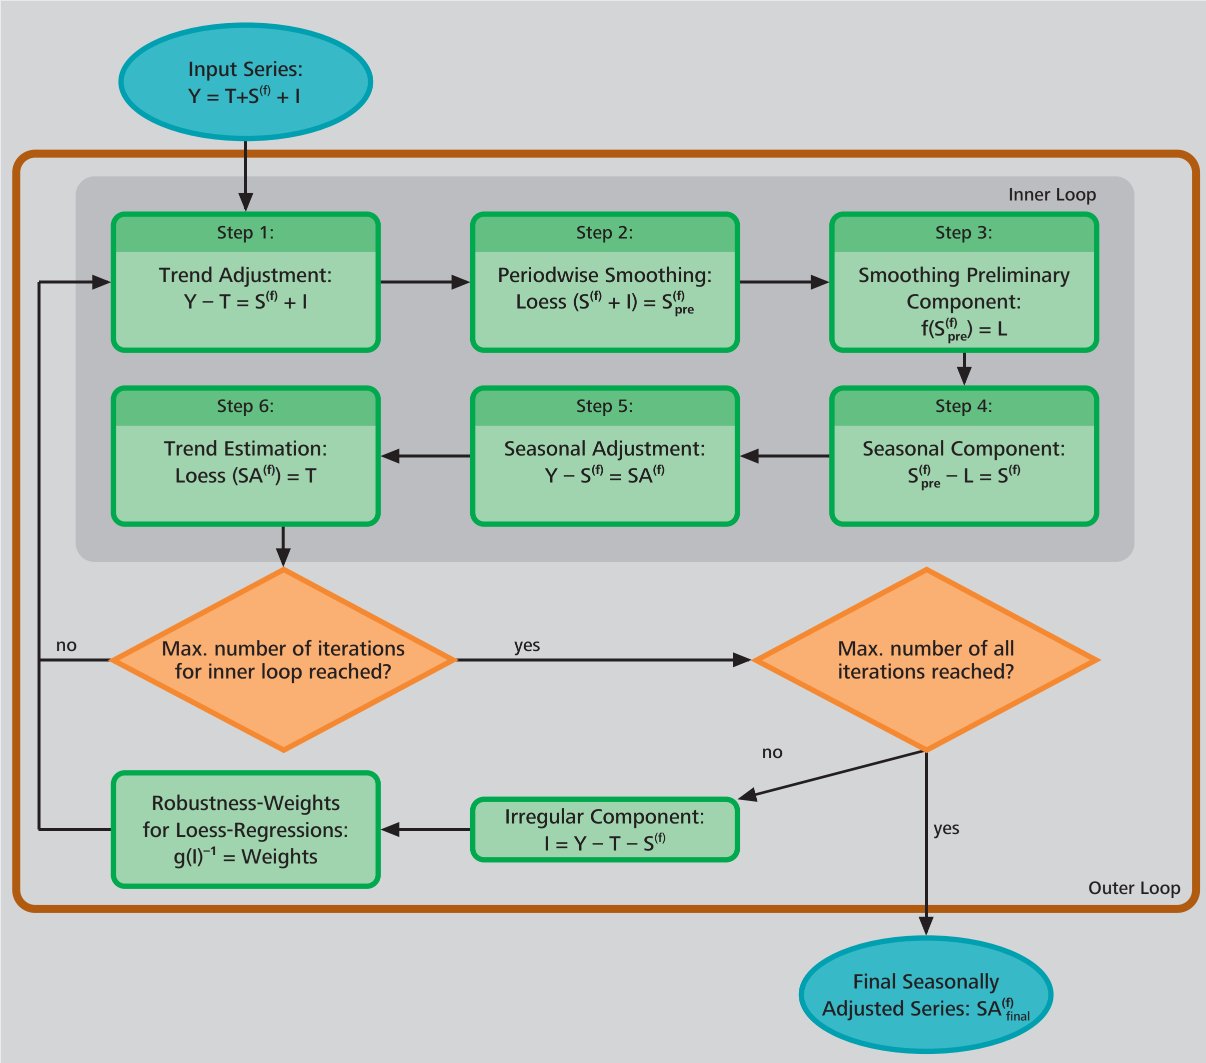

Deutsche Bundesbank

3 Daily Seasonal Adjustment

Daily time series can contain three different seasonalities. This can be illustrated with the daily time series of the currency in circulation[2]: if people primarily need money in cash to buy groceries for the weekend, the data will display a weekly pattern of increasing currency in circulation towards the end of the week. An increase around the end of the month of the circulating currency results from an increased withdrawal of cash following the payment of salaries. Finally, consumption increases due to Christmas, and thus, the

[2]For a comprehensive discussion of the factors determining supply and demand of the currency in circulation see Bartzsch, Seitz, and Setzer (2015) and Cabrero, Camba-Mendez, Hirsch, and Nieto (2009).

currency in circulation increases at the end of each year. Correspondingly, many daily time series are characterised by weekly, monthly and annually recurring patterns. In addition, travel activities around moving holidays might increase the demand for banknotes as well. Thus, the model described in Equation 1 has to be refined to incorporate intra-weekly, intra-monthly and intra-annual seasonality, $S_t^{(7)}$, $S_t^{(31)}$, and $S_t^{(365)}$, as well as moving holiday effects C_t:

$$Y_t = T_t + S_t^{(7)} + S_t^{(31)} + S_t^{(365)} + C_t + I_t. \tag{8}$$

STL adjusts each seasonal factor separately, and therefore has to be run three times in order to capture all $S_t^{(f)}$ $\forall\, f \in (7, 31, 365)$.

The idea of iteratively adjusting several seasonal frequencies has already been put forward by Cleveland et al. (1990). In the case of daily time series, the main obstacle is the constraint of STL to have a constant period length, i.e. the number of observations has to be identical in each period. This is not unique to STL, though. For example, seasonal RegARIMA models usually face the same requirements. While this is unproblematic in the case of intra-weekly seasonality, the number of days per month and the number of days per year are not fixed. As a result, the period length has to be standardised, either by omitting a subset of the data or by artificial prolongation. Possible implementations are discussed below.

A further task is the development of a calendar and outlier adjustment routine together with the forecasting of seasonal factors. The latter is necessary for data producers providing seasonally adjusted figures applying a current adjustment revision policy, i.e. the policy of reidentifying time series components at set review periods and adjusting new data points by forecasted seasonal factors in the meantime (Eurostat, 2015). Outlier adjustment can be important for an unbiased estimation of the calendar effects as well as for the prediction of the future trajectory of the time series. The calendar and outlier adjustment needs to be integrated into the iterative seasonal adjustment.

As in X-13 and TRAMO/SEATS, the RegARIMA model could be the basis for all prior adjustments including estimating calendar effects. In the case of daily time series, several arguments can be made for changing this order and to begin with adjusting the intra-weekly movements: firstly, each year of data contains more than 50 observations of every weekday. The adjustment of the intra-weekly seasonality therefore tends to be robust against outliers and rare effects such as moving holidays even in the case of short series. Secondly, the intra-weekly, intra-monthly and intra-annual periodic movements necessitate different data modifications that are somewhat antagonistic. While the weekday effects need the full set of available data, the adjustment of the effect of the day of the year entails the standardisation of the length of the year to be either 365 or 366. If we started by estimating the effects of the moving holidays, there would be a risk that the estimated effects contain some of the weekday pattern.

Taking all of these considerations into account results in the following algorithm:

- Step I: Adjusting Intra-Weekly Seasonality with STL.
- Step II: Calendar- and Outlier Adjustment with RegArima.
- Step III: Adjusting Intra-Monthly Seasonality with STL.
- Step IV: Adjusting Intra-Annual Seasonality with STL.

It can of course be the case that a given time series does not contain all of the seasonal and calendar components. In that case it will be advisable to omit the respective seasonal adjustment steps or leave out regressors that capture the effects of moving holidays.

3.1 Step I: Adjusting Intra-Weekly Seasonality

Before estimating the intra-weekly seasonality with STL, the series may have to be pre-adjusted. This may include the interpolation of missing values or taking the logarithm of the data to stabilise the variance of the series.

For most of the computations in STL there exist sensible default values for the smoothing parameters (see Cleveland et al., 1990). Only for the smoothing of the period-wise sub-series $\gamma_{S^{(7)}}$, i.e. the number of observations included in the local regression to obtain the seasonal factor, a specification is needed. In the case of intra-weekly seasonality, $\gamma_{S^{(7)}}$ has to be large enough so that the weekday effects do not get confounded by other effects such as moving holidays. Obviously, the parameter $\gamma_{S^{(7)}}$ should not be set too high either in order to capture changes in the seasonal pattern. As the time series has not been outlier adjusted yet, usually, the robust version of STL is preferable, so the number of iterations of the outer loop, which can be set by the user, is one or higher.

3.2 Step II: Calculating the Effects of Moving Holidays and Outliers

Following the notation of Ghysels and Osborn (2001) a seasonal RegARIMA model for a time series with 365 observations per year can be written as

$$\phi_p(B)\phi_P(B^{365})(1-B)^d(1-B^{365})^D\left(y_t - \sum_{i=1}^{r}\beta_i x_{it}\right) = \theta_q(B)\theta_Q(B^{365})\varepsilon_t \tag{9}$$

where $\phi(B)$ and $\theta(B)$ are polynomials of order p, P, q, and Q, while B is the backshift operator. β_i captures the impact of the i-th regressor x_{it} on the time series y_t and ε_t is the error term. The ARIMA part of Equation 9 can be written in short form as $(p\,d\,q)_1(P\,D\,Q)_{365}$ with the AR order p, MA order q and number of differences d and capitals indicating seasonal terms (for a thorough discussion of ARIMA modelling see Box and Jenkins, 1970; Brockwell and Davis, 1987).

As a higher periodicity increases the computational burden, it is not always feasible to use seasonal ARIMA models with daily series. In these situations, popular alternatives include estimating the seasonal factor by trigonometric functions or dummy variables in a non-seasonal RegARIMA model. The former can be expressed as a function of sines and cosines (Campbell and Diebold, 2005; De Livera, Hyndman, and Snyder, 2011):

$$\hat{S}_t = \sum_{j=1}^{J} \left(\hat{\beta}_{1,j}\, sin\left(\frac{2\pi j G(t)}{365}\right) + \hat{\beta}_{2,j}\, cos\left(\frac{2\pi j G(t)}{365}\right)\right) \tag{10}$$

where $G(t)$ is a unit step function indicating the day of the year, i.e. cycles through 1, ..., 365. The optimal size of J is determined by the corrected Akaike information criterion (AICc).[3] J will usually be no larger than 30 and thus the trigonometric variant is more parsimonious than the dummy variable approach. It is therefore the preferred option in the DSA procedure.

The order of the ARIMA model is either specified by the user or can be determined by automatic model detection. For the latter the Hyndman-Khandakar algorithm is implemented (Hyndman and Khandakar, 2008), though other methods could be used as well (e.g. Gómez and Maravall, 2001).

The outlier adjustment is similar to that described by Chen and Liu (1993).[4] The procedure consists of two parts. Part one contains the iterative one-by-one detection of possible outliers at all time points based on the residuals of an ARIMA model. Part two includes all potential outliers as regressors in a final ARIMA model and then removes all candidates from the set of potential outliers that do not surpass a prespecified threshold for the t-statistics. For monthly time series, these two parts are iterated until convergence. In the case of daily time series, convergence is infrequent and divergence, i.e. an increasing number of outliers found in each iteration, a serious possibility. To avoid this, the threshold for keeping outliers is higher, often a critical value of the t-statistic of 7 or

[3]In our case, the AICc is preferable to the Bayesian information criterion (BIC) as it asymptotically selects the model that minimises the mean squared prediction error (Vrieze, 2012). Also, empirical results show that sometimes the BIC selects a value for J that is close or equal to one, so that only basic annual and no monthly movements are estimated, even if intra-monthly seasonality is clearly present.

[4]That procedure is implemented in the R package tsoutliers (López, 2015). The procedure described here adapts some of the implemented methods to the features of daily time series.

more is used. Additionally, the number of iterations may be chosen to be smaller than for low frequency time series.

If one or several moving holidays have an impact on the original time series, regressors for these holidays can be incorporated into the estimation process. Often, it can be difficult to decide on an optimal functional form of the regressors, which always will depend on the characteristics of the given time series. Section 4.1 discusses alternatives for the currency in circulation. Evidently, the impact of other variables could be accounted for by including appropriate regressors into the RegARIMA model.

For data producers carrying out seasonal adjustment on a regular basis, it can become time-consuming to recalculate the complete seasonal adjustment every time a new data point is added to a series. To overcome the necessity of daily re-estimation, forecasted seasonal factors can be used for the adjustment of additional data. We can extend the series adjusted for intra-weekly seasonality using forecasts of length h based on the final RegARIMA model. As the next two steps, i.e. removing intra-monthly and intra-annual seasonality, are based on the series adjusted for intra-weekly and moving holiday effects, the final seasonally adjusted series will include an h-step ahead forecast. Thus, it is straightforward to compute forecasts for the intra-monthly and intra-annual seasonal factors. The forecasts of the intra-weekly seasonal factors have to be obtained independently.

The forecast of the original series will be given by the forecast of the seasonal factor of the weekday, $S_t^{(7)}$, and the forecast of the original series adjusted for weekday effects, $SA_t^{(7)}$:

$$\hat{Y}_{t+h} = \hat{S}_{t+h}^{(7)} + \widehat{SA}_{t+h}^{(7)}. \tag{11}$$

As discussed, $\widehat{SA}_{t+h}^{(7)}$ is obtained using the RegARIMA model. $\hat{S}_{t+h}^{(7)}$ is forecasted using Holt's (2004) double exponential smoothing with heterogenous seasonal effects to extrapolate the seasonal factors for the day of the week.[5] The level l_t and the slope b_t of $S_t^{(7)}$ are given by:

$$l_{tj} = \alpha_j S_{tj}^{(7)} + (1-\alpha_j)(l_{t-7,j} + b_{t-7,j}) \qquad \forall \ j \in \ (1, ..., 7) \tag{12}$$

$$b_{tj} = \mu_j (l_{tj} - l_{t-7,j}) + (1-\mu_j) b_{t-7,j} \qquad \forall \ j \in \ (1, ..., 7). \tag{13}$$

with α_j and μ_j estimated from the data. This is equivalent to estimating double expo-

[5] In the Census X-11 method without RegARIMA modelling, the seasonal factors are forecasted assuming that the change of each period's seasonal factor is proportional to the last year's change (Ladiray and Quenneville, 2001). For a given number of observations per seasonal cycle f, the seasonal factors are extrapolated by $S_{t+f} = S_t + \alpha^\star (S_t + S_{t-f}) \ \forall \, t \in (N+1-f, ..., N+h-f)$ with $\alpha^\star$ usually set to 1/2. The disadvantage of this heuristic is that the rate of change $\alpha^\star$ has to be set by the user. Also, for daily data, h can become quite high. This increases the importance of estimating rather than specifying $\alpha^\star$.

nential smoothing equations for each weekday separately. Equations 12 and 13 are used to obtain an h-step ahead forecast:

$$\hat{S}^{(7)}_{t+h,j} = l_{tj} + hb_{tj}. \tag{14}$$

Adding the forecast of the intra-weekly seasonal factors to the RegARIMA forecasts yields the forecasts for the original series. These can be used to compute implicit seasonal factors given by the difference between the original and the seasonally adjusted time series.

3.3 Step III: Adjusting Intra-Monthly Seasonality

While the length of any year is always the same except for the leap year, the number of days in a month ranges 28 to 31 with an average of 30.4 days per month. Ignoring the change of the number of days by assuming each month to have 30 days would blend together the effects of different days of the month.

The problem of confused effects can be circumvented by extending each month to 31 days. There are basically two methods of achieving this. Either the course of each month has to be stretched to 31 days by interpolation after warping the time axis, or the deficit days have to be filled up. The characteristics of the intra-monthly seasonality determine which of these options is preferable. If the seasonal pattern is compressed in shorter months, i.e. peaks and troughs have the same magnitude but occur in less time, then stretching is the optimal response. The latter method is preferable if the typical seasonal movement of months with 31 days is cut-off or halted in the other months.

For the extension of the time series, several different methods can be utilised. While regression based techniques have the advantage of possibly incorporating much or all of the information that is contained in the actual series, in the context of time series analysis they can become computationally expensive. Therefore, it seems preferable to use a more parsimonious approach, namely cubic splines that can be used for stretching or extending the time series to span over 31 days.[6] Efficient algorithms are readily available and have the benefit of a high degree of smoothness, and thus are an alternative to simple linear interpolation that does not take into account the local dynamics of the series, though differences are usually small. The smoothness is ensured as cubic splines have continuous first and second derivatives by definition.

For the i^{th} interval that ranges from t_i to t_{i+1}, the spline function for our interpolated time series $Y^\star$ can be represented by a simple cubic polynomial

$$Y_i^\star(t) = a_{1,i}(t-t_i)^3 + a_{2,i}(t-t_i)^2 + a_{3,i}(t-t_i) + a_{4,i}. \tag{15}$$

[6] Depending on the characteristics of the time series other interpolation techniques may be preferred. If smoothness is not desirable, linear interpolation or carrying the last observation forward can be used.

To obtain a solution for the parameters $a_{1,i}, a_{2,i}, a_{3,i}$, and $a_{4,i}$ the first and second derivatives are required to be continuous, i.e.

$$\frac{dY_i^\star(t_i)}{dt} = \frac{dY_{i-1}^\star(t_i)}{dt} \tag{16}$$

and

$$\frac{d^2Y_i^\star(t_i)}{dt^2} = \frac{d^2Y_{i-1}^\star(t_i)}{dt^2} := \eta, \tag{17}$$

where η has to be specified in order to obtain a unique solution (Forsythe, Malcolm, and Moler, 1977; Press, Teukolsky, Vetterling, and Flannery, 1992). Here, the Forsythe-Malcolm-Moler algorithm is applied to obtain values for the additional data points needed to have 31 days in each month. Then, the intra-monthly seasonal influences are estimated based on the resulting time series, for which $\gamma_{S^{(31)}}$ has to be chosen to fit the characteristics of the intra-monthly seasonal pattern.

3.4 Step IV: Adjusting Intra-Annual Seasonality

As the time series has been extended in step II, the excess days including every 29th of February are removed so that each year contains 365 days. Then $\gamma_{S^{(365)}}$ has to be chosen: now that potential outliers and the influence of moving holidays have been removed from the series, the main objective in choosing $\gamma_{S^{(365)}}$ is capturing the variability of the intra-annual seasonality optimally. As in step I and III, the intra-annual seasonality is estimated and adjusted for using STL.

3.5 Post-adjustment and Forecast

As discussed in Section 3.2, the forecasts of the original series are calculated as the sum of the ARIMA based forecasts of the original series adjusted for intra-weekly seasonality and the forecasts of the intra-weekly seasonal factors. The main objective of these forecasts is the provision of seasonal factors so that data providers can use a (controlled) current adjustment scheme for the adjustment of their data (Eurostat, 2015). This is of practical importance, because it is time-consuming to re-identify the model, the filters, and the regression parameters each time new data are available.

After all seasonal factors have been estimated, the 29th of February is added to the seasonally adjusted series using a spline interpolation. Finally, the effects of the outliers are added to the seasonally adjusted time series.

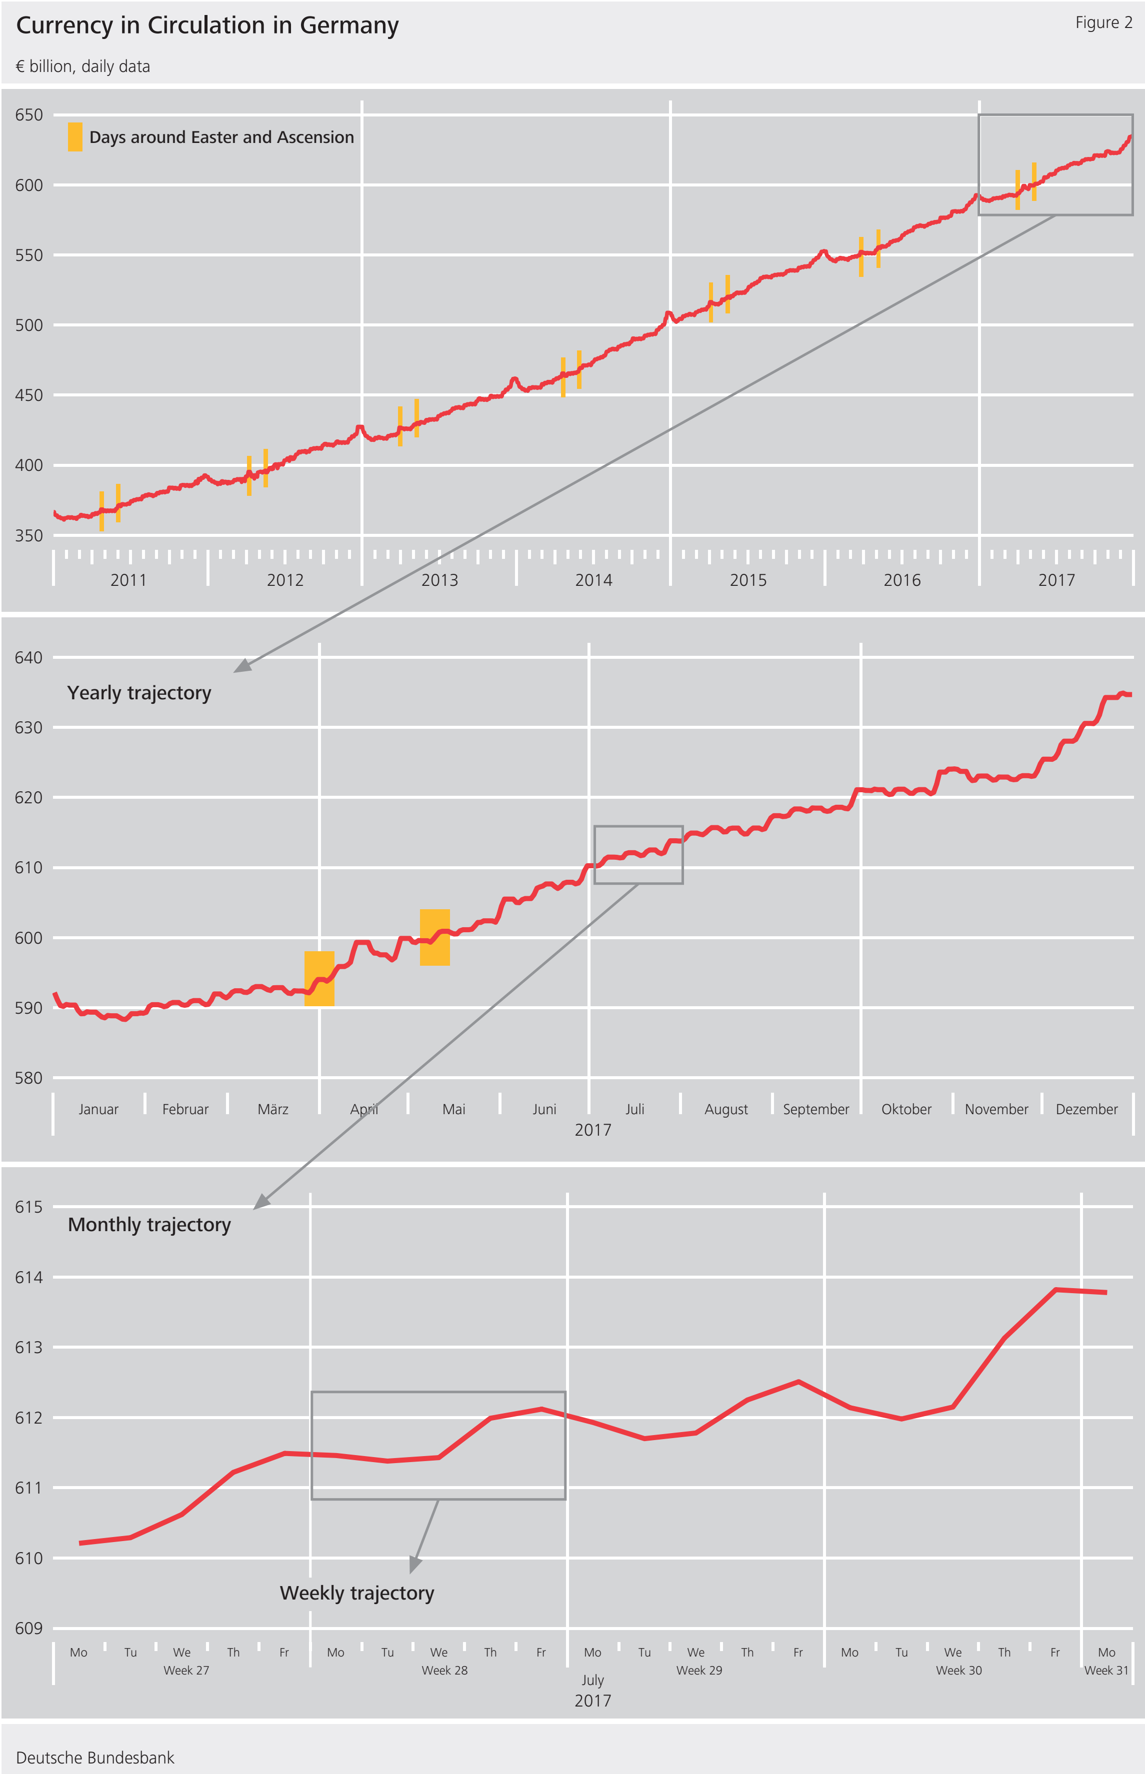
Currency in Circulation in Germany
Figure 2
€ billion, daily data
Days around Easter and Ascension
650
600
550
500
450
400
350
2011
2012
2013
2014
2015
2016
2017
Yearly trajectory
640
630
620
610
600
590
580
Januar
Februar
März
April
Mai
Juni
Juli
August
September
Oktober
November
Dezember
2017
Monthly trajectory
615
614
613
612
611
610
609
Weekly trajectory
Mo
Tu
We
Th
Fr
Week 27
Week 28
Week 29
Week 30
Week 31
July
2017
Deutsche Bundesbank

4 Application: Adjusting the German Currency in Circulation

The currency in circulation is calculated as the cumulated differences of cash outflows from and inflows to the central bank. The distribution of banknotes to the public runs primarily through commercial banks. The series' trend is closely related to nominal GDP, but is also marked by periodical influences (Cabrero et al., 2009). Figure 2 highlights the different seasonal fluctuations of the time series. Most noticeably, the currency in circulation increases in Germany some weeks before Christmas, peaking just before the end of December. In addition, the series is characterised by an increase towards the end of most months as well as towards the end of the week. Finally, an increased demand for banknotes can be observed around moving holidays such as Easter and Pentecost.

For the currency in circulation in Germany data for all denominations from 5 to 500 Euro are available on a daily basis except for the weekend. The time series contain about 7 years of data, starting in 2011. As no new currency is circulated on the weekend it is plausible to assume that the number of banknotes is the same as on the previous Friday. Here, the discussion will focus primarily on the total currency in circulation, but results for each denomination are shown in Appendix C.

4.1 Specifications of the Seasonal and Moving Holiday Adjustment

Using the robust STL, the value for $\gamma_{S^{(7)}}$ that controls the stability of the estimated seasonal factor is set to 151, i.e. about 3 years of neighbouring observations are included in the local regression. This is sufficiently high to avoid including annual periodic movements into the intra-weekly seasonal factors, while also taking into account the high stability of the weekly seasonal pattern. To assess the flexibility of the seasonal factors, one strategy is to inspect the SI-ratio, i.e. the combined seasonal-irregular component, together with the seasonal factor of interest. Both should follow a similar trajectory, except that the seasonal-irregular component will fluctuate more (see Cleveland and Terpenning, 1982).

As already mentioned, it will often prove difficult to find the optimal functional form of the regressors that capture the influence of moving holidays and which holidays to include. In the case of the currency in circulation, it seems plausible to assume that Easter and Ascension have an effect on the demand for banknotes. This is because more money is spent, for example on presents, travel or other activities. Whether or not people decide to travel during their holidays depends in part on the dates of Easter and Ascension. In addition, the exact timing of an increase in money demand depends on random events that change from year to year. The effects of these contingencies tend to even out over the course of a month. But for daily time series, the objective is to estimate the influence of Easter on the currency in circulation of every potential day and thus, the available data can be quite noisy.

One option for modelling calendar effects is to use a dummy variable for each of m days before and for each of n days after Easter Sunday and for the other holidays resulting in $m + n + 1$ dummy variables per holiday. Obviously, this approach has its drawbacks. Firstly, we face a potentially large loss of degrees of freedom. Secondly, the model may be overfitting, especially if only a few years of observations are available. In that case, the estimate of the effect of at least some regressors will lead to a poor out-of-sample fit. And lastly, m and n have to be chosen. Using information criteria such as the AICc the number of days to be included before and after Easter can be determined. As this means comparing a large set of different models, a more practical approach is to estimate a RegARIMA model with large values of m and n and then delete all non-significant dummy variables.[7]

The number of cosine and sine terms to be included in Equation 10 to capture the monthly and annual seasonal movements is also determined by the lowest AICc which in this case leads to the inclusion of 28 x 2 trigonometric terms. This suggests that intra-annual as well as intra-monthly movements have a significant influence. Note here that multiples of 12 capture monthly patterns. The RegARIMA results show that the effects of Easter and Ascension are significant and 13 dummy variables are needed as $m = 7$ and $n = 5$.

The outlier detection found three outliers that all impact the series in the end of 2012 or beginning of 2013.

For the computation of the intra-monthly seasonal factors the length of the seasonal Loess filter $\gamma_{S(31)} := 51$, i.e. more than 4 years of observations are included in the estimation of each value of the intra-monthly seasonal factor. Additionally, the robust version of STL is used. This ensures that annual effects do not influence the computations. As the series adjusted in this step includes the forecasts from the RegARIMA models, the intra-monthly seasonal factors already include forecasted values.

For the final step, choosing a long seasonal filter results in an almost stable seasonal factor, i.e. the value of the intra-annual seasonal factor is close to being identical in all years. On the other hand, a too small value for $\gamma_{S(365)}$ may lead to excessive variability

[7] An alternative that uses fewer regressors can be derived by noting that usually there is an increase in the currency in circulation before Easter Sunday that fades out but lasts until a few days after. A similar relation can be found for Ascension, although less pronounced. For Easter, this pyramid-like pattern can be modelled parsimoniously by means of the function $x_t = (\frac{n+m+2}{2} - |i - \frac{-m+n}{2}|)\mathbb{1}_{t=Easter+i} \ \forall \ i \in [-m; n]$.

Log Spectrum of the Currency in Circulation

Differenced series

Figure 3

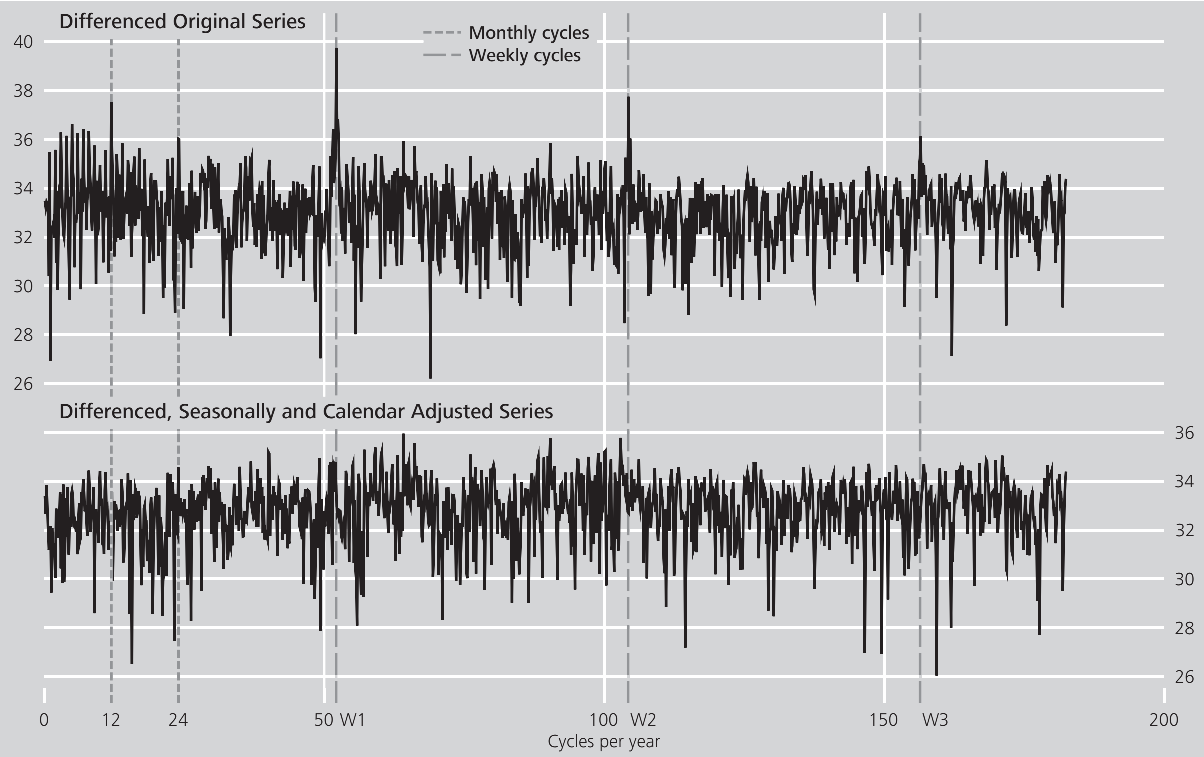

Deutsche Bundesbank

in the trajectory of the seasonal factor. To avoid single observations dominating the estimation of the local regressions and moving averages, again, the robust version of STL is used, the number of inner and outer loops of the STL algorithm is increased, and the Loess parameter for the seasonal factor $\gamma_{S(365)} := 13$.

Comparing the smoothed periodogram (Venables and Ripley, 2002) of the differenced original and the differenced final adjusted series indicates that the spectral peaks at the seasonal frequencies have been filtered out (see Figure 3. Frequency 12 and 24 are the first monthly cycles; W1, W2, W3 are the first three weekly cycles corresponding to $\frac{365}{7}$, $\frac{730}{7}$, $\frac{1095}{7}$ cycles a year). So based on the spectral diagnostics, we can conclude that the seasonal movements of the currency in circulation have been eliminated successfully. Figure 4 shows the original series of the total currency in circulation in Germany and its calendar and seasonally adjusted counterpart.

Seasonal Adjustment Results for the Currency in Circulation in Germany

Figure 4

in Billion

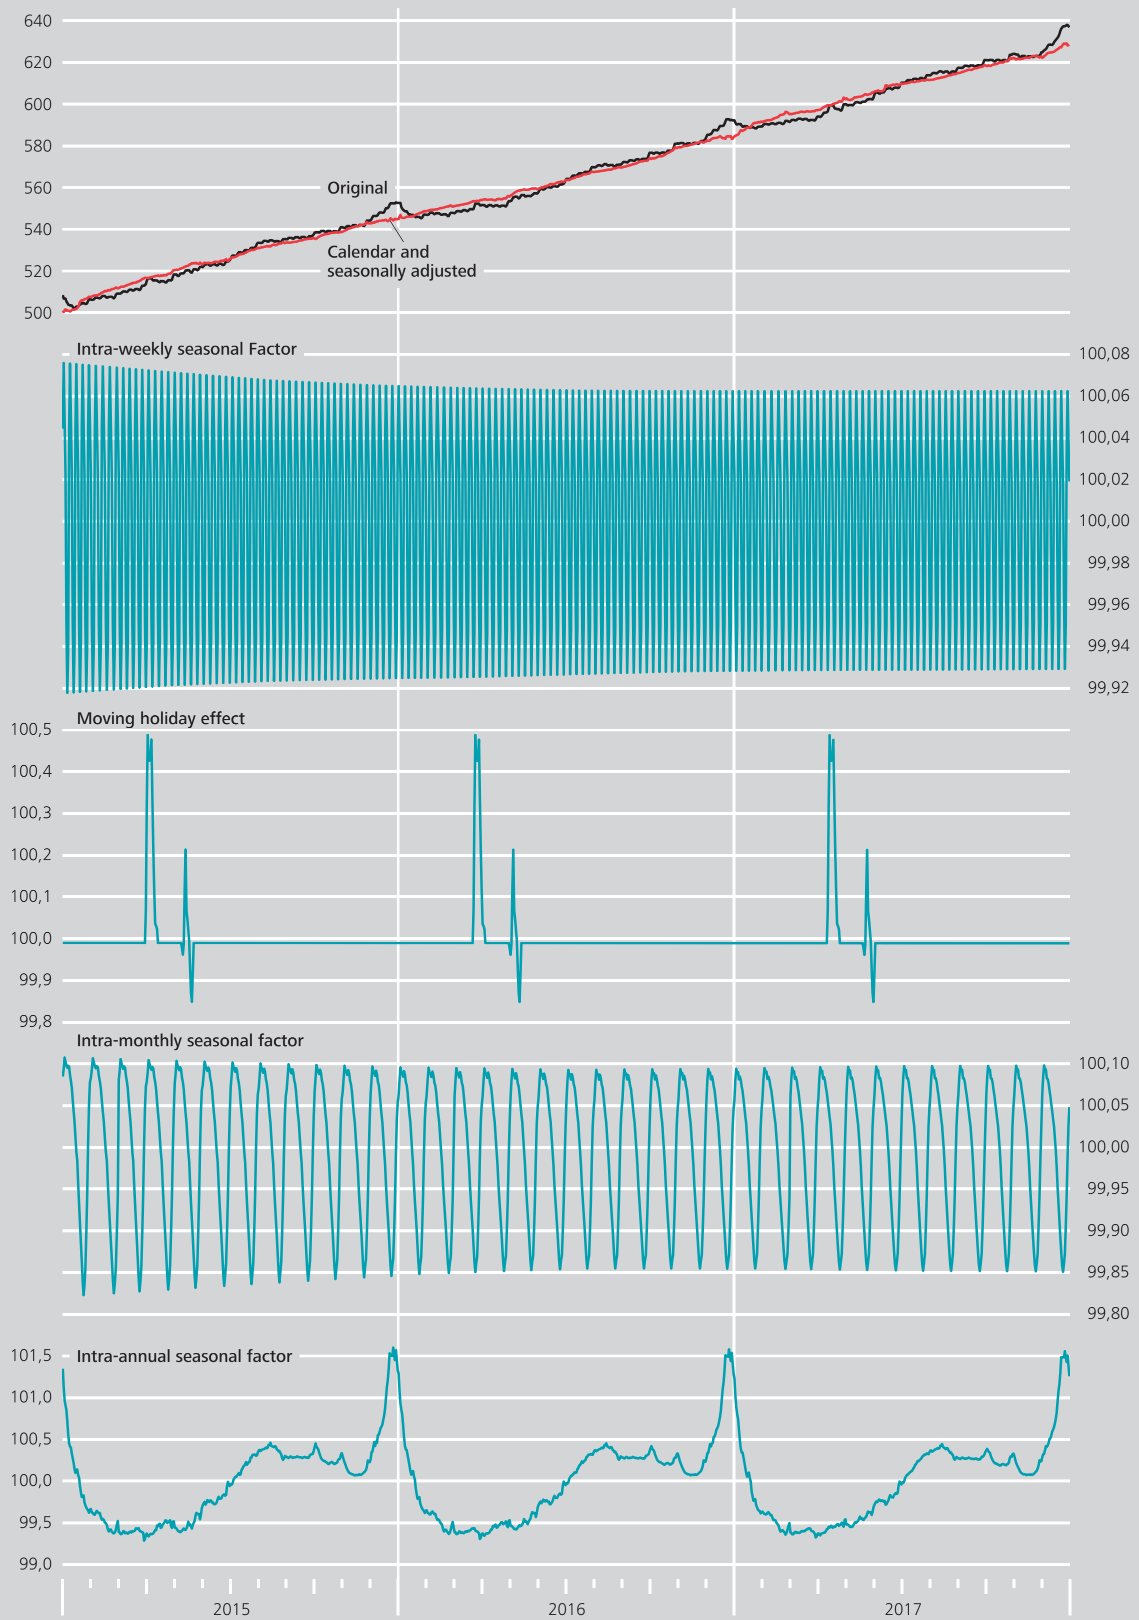

Deutsche Bundesbank

Table 1: Mean Absolute Deviation (MAD), Root Mean Square Deviation (RMSD, Both in Million Euro), and Mean Absolute Percentage Deviation (MAPD) of the Seasonally Adjusted Currency in Circulation in Germany Obtained Using DSA from Results Obtained by Other Methods.

		Time series aggregated to					
		monthly means			month-end values		
		MAD	RMSD	MAPD	MAD	RMSD	MAPD
5 Euro	TRAMO/SEATS	19.51	22.80	0.23	22.06	25.70	0.26
	X-13	20.32	22.99	0.24	22.03	25.58	0.26
	STR	22.15	26.16	0.26	23.02	27.98	0.27
	STR*	23.76	28.37	0.28			
10 Euro	TRAMO/SEATS	24.95	32.14	0.11	62.80	73.68	0.27
	X-13	23.20	29.34	0.10	64.30	74.08	0.28
	STR	32.33	43.39	0.15	48.31	66.94	0.22
	STR*	40.80	57.76	0.19			
20 Euro	TRAMO/SEATS	63.33	79.59	0.17	160.81	189.03	0.44
	X-13	67.26	83.93	0.18	153.35	185.09	0.43
	STR	96.37	118.44	0.27	111.16	139.81	0.30
	STR*	106.57	133.17	0.29			
50 Euro	TRAMO/SEATS	201.55	240.02	0.14	384.08	455.45	0.27
	X-13	200.83	242.94	0.14	431.69	492.68	0.31
	STR	210.19	254.55	0.14	324.39	417.41	0.23
	STR*	263.00	334.68	0.18			
100 Euro	TRAMO/SEATS	95.10	116.48	0.11	109.76	141.30	0.13
	X-13	128.91	165.13	0.13	131.08	156.33	0.14
	STR	240.11	296.46	0.25	248.41	315.93	0.26
	STR*	247.22	310.25	0.26			
200 Euro	TRAMO/SEATS	41.38	58.51	0.18	49.70	59.96	0.20
	X-13	38.73	53.62	0.15	39.95	52.32	0.15
	STR	43.65	55.94	0.17	43.34	56.47	0.17
	STR*	48.81	68.43	0.20			
500 Euro	TRAMO/SEATS	156.46	217.31	0.11	171.28	226.84	0.12
	X-13	175.70	234.28	0.11	197.38	251.97	0.13
	STR	420.10	574.59	0.27	439.81	591.26	0.29
	STR*	436.53	596.60	0.28			
Total Euro	TRAMO/SEATS	326.11	390.54	0.07	728.78	850.39	0.15
	X-13	353.96	437.62	0.07	762.06	867.99	0.16
	STR	429.14	525.70	0.09	537.71	692.39	0.11
	STR*	498.21	634.72	0.11			

* STR without any temporal aggregations, i.e. using daily figures.

4.2 Comparison with other seasonal adjustment procedures

We compare the results to those obtained from X-13 and TRAMO/SEATS in the JDemetra+ package (version 2.2) by transforming the daily figures to end-of-month values and monthly means respectively (see Figure 6 and 7 in Appendix C). Additionally, we include results obtained from using the seasonal adjustment routine STR developed by Dokumentov and Hyndman (2015).

For the time series aggregated to monthly means, neither the RegARIMA procedure in X-13 nor the TRAMO method in TRAMO/SEATS automatically detects Easter effects. If we include these manually the regressors are not significant. They are though, if daily data are used in the DSA procedure (see Table 3 in Section B). Of course, these results depend on how the Easter effect is modelled. The built-in regressors in JDemetra+ indicate the location of the six respectively eight days before Easter, while a slightly different variant is used for the daily time series (see Section 3.2). Similarly, effects of the number of weekdays have not been detected by the common seasonal adjustment procedures. The availability of more data and equivalently the higher "visibility" of weekly patterns is one major advantage of the seasonal adjustment of daily compared to lower frequency time series.

Overall, for monthly aggregates, the discrepancies between the DSA procedure, X-13 and TRAMO/SEATS are fairly small for the currency in circulation (see Table 1). The mean absolute percentage difference (MAPD) between the DSA result and those obtained by either TRAMO/SEATS or X-13 ranges from 0.08 to 0.24 percent for all series aggregated to monthly means. Correspondingly, the MAPD for these series ranges from 0.07 to 0.24 percent for the implicit combined calendar and seasonal factor. For the series aggregated to end-of-month values the MAPD ranges from 0.12 to 0.44 percent for the seasonally adjusted series as well as for the implicit factors. The difference to STR is similar and depending on the data transformation used, the difference is even smaller. Measured by MAPD the differences between DSA and STR range from 0.09 to 0.27 percent based on monthly averages and 0.11 to 0.30 percent based on end-of-month values. If we use the daily figures without any temporal aggregation, the MAPD lies between 0.11 and 0.29 percent.

The small differences may be a consequence of the structure of the time series, especially given that the seasonal influences are dominated by a strong trend. Therefore, as a further comparison, a small set of daily time series of different length was simulated by means of a structural model, where an ARIMA process is combined with seasonal factors,

Table 2: Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Mean Absolute Percentage Error (MAPE) of the Estimated Seasonal Factors

	Simulated daily time series aggregated to monthly means			month-end values		
	MAE	RMSE	MAPE	MAE	RMSE	MAPE
DSA	2.18	2.69	2.83	2.51	3.15	3.21
TRAMO/SEATS	2.71	3.37	3.51	3.34	4.25	4.38
X-13	2.17	2.74	2.72	3.13	3.89	4.10
STR	2.37	3.01	3.23	2.61	3.30	3.55

$S_t^{(f)}\ \forall\ f \in (7, 31, 365)$, obtained from trigonometric terms and time-varying effects:

$$Y_t = SA_t + S_t^{(7)} + S_t^{(31)} + S_t^{(365)}. \tag{18}$$

$$SA_t = \frac{(1-\theta_1)}{(1-\phi_1 B - \phi_2 B^2 - \phi_3 B^3)(1-B)}\varepsilon_t. \tag{19}$$

$$S_t^{(f)} = \sum_{j=1}^{J^{(f)}} \left(\beta_{j,t}^{(f)} sin\left(\frac{2\pi j G(t)}{f}\right) + \beta_{j,t}^{(f)} cos\left(\frac{2\pi j G(t)}{f}\right)\right). \tag{20}$$

$$\beta_{j,t}^{(f)} = \rho^{(f)} \beta_{j,t-1}^{(f)}. \tag{21}$$

The time series are simulated with $\theta_1 := -0.4$, $\phi_1 := -0.2$, $\phi_2 := 0.5$, $\phi_3 := 0.1$, error terms are drawn from a standard normal distribution, for the weekly seasonal pattern $\rho^{(7)} \sim \mathcal{N}(1, 10^{-4})$ and $\beta_{j,0}^{(7)} = 1.6 \cdot 0.7^j$, for the monthly seasonal pattern $\rho^{(31)} \sim \mathcal{N}(1, 0.00015)$ and $\beta_{j,0}^{(31)} = 1.6 \cdot 0.6^j$, and for the annual seasonal pattern $\rho^{(365)} \sim \mathcal{N}(1, 0.00025)$ and $\beta_{j,0}^{(365)} = 4.4 \cdot 0.9^j$.

Visual inspection of Figure 8 and 9 in the Appendix corroborates the assertion from before that DSA produces seasonally adjusted series similar to those obtained by X-13 and TRAMO/SEATS.

For the series obtained through taking means over the whole month, the mean absolute error (MAE) of the estimated seasonal factors are 2.2, 2.7, 2.2, and 2.4 for DSA, TRAMO/SEATS, X-13, and STR. Based on the end-of-the-month figures, the MAEs are 2.5, 3.3, 3.1, and 2.6. For the seasonally adjusted series without any temporal aggregation, the MAE is 2.5 and 2.6 for DSA and STR, respectively (see also Figure 10). Thus, the MAE do not find that any method completely outperforms all others, as somewhat similar seasonally adjusted time series are obtained. The same conclusions can be drawn from the root mean squared error (RMSE) and the mean absolute percentage error (MAPE) as can be seen in Table 2. Again, this indicates a high validity of the estimated seasonally adjusted series.

5 Discussion

The procedure presented here closes a gap in the methods for seasonal and moving holiday adjustment of daily time series. As shown with the currency in circulation and a set of simulated series, the procedure yields sensible results and is able to adjust the intra-weekly, intra-monthly and intra-annual periodic movements in addition to moving holidays. The comparison of the aggregated daily series with the results obtained by using the well-established seasonal adjustment procedures for monthly time series implemented in JDemetra+ indicates the validity of the results of DSA.

Future developments should aim at improving the computational speed and reliability of the outlier detection and estimation. Also, optimal default settings for the length of the seasonal Loess filter $\gamma_{S(f)}$ should be investigated. Post-smoothing of the seasonal factors may increase the stability of the seasonal factors for series with highly harmonic, steady seasonal movements and few observations. Finally, reliable seasonality tests that are attuned to the characteristics of daily time series are needed.

References

Akaike, H. (1980). Seasonal Adjustment by a Bayesian Modeling. *Journal of Time Series Analysis 1*(1), 1–13.

Akaike, H. and M. Ishiguro (1983). Comparative study of the x-ll and baysea procedures of seasonal adjustment. In A. Zellner (Ed.), *Applied Time Series Analysis of Economic Data*, pp. 17–50. New York, US.

Bartzsch, N., F. Seitz, and R. Setzer (2015). The Demand for Euro Banknotes Issued in Germany: Structural Modelling and Forecasting. Technical report, ROME Discussion Paper Series.

Box, G. and G. Jenkins (1970). *Time Series Analysis: Forecasting and Control.* San Francisco: Holden-Day.

Brockwell, P. and R. Davis (1987). *Time Series: Theory and Methods.* New York: Springer.

Cabrero, A., G. Camba-Mendez, A. Hirsch, and F. Nieto (2009). Modelling the Daily Banknotes in Circulation in the Context of the Liquidity Management of the European Central Bank. *Journal of Forecasting 28*(3), 194–217.

Campbell, S. D. and F. X. Diebold (2005). Weather Forecasting for Weather Derivatives. *Journal of the American Statistical Association 100*(469), 6–16.

Chen, C. and L.-M. Liu (1993). Joint Estimation of Model Parameters and Outlier Effects in Time Series. *Journal of the American Statistical Association 88*(421), 284–297.

Cleveland, R. B., W. S. Cleveland, J. E. McRae, and I. Terpenning (1990). STL: A Seasonal-Trend Decomposition Procedure Based on Loess. *Journal of Official Statistics 6*(1), 3–73.

Cleveland, W. S. and S. J. Devlin (1988). Locally Weighted Regression: An Approach to Regression Analysis by Local Fitting. *Journal of the American statistical association 83*(403), 596–610.

Cleveland, W. S. and I. J. Terpenning (1982). Graphical Methods for Seasonal Adjustment. *Journal of the American Statistical Association 77*(377), 52–62.

De Livera, A. M., R. J. Hyndman, and R. D. Snyder (2011). Forecasting Time Series with Complex Seasonal Patterns Using Exponential Smoothing. *Journal of the American Statistical Association 106*(496), 1513–1527.

Dokumentov, A. and R. J. Hyndman (2015). STR: A Seasonal-Trend Decomposition Procedure Based on Regression. Technical report, Monash University, Department of Econometrics and Business Statistics.

Eurostat (2015). *ESS Guidelines on Seasonal Adjustment.*

Forsythe, G. E., M. A. Malcolm, and C. B. Moler (1977). *Computer Methods for Mathematical Computations.* Prentice-Hall Series in Automatic Computation. Englewood Cliffs, New Jersey. Prentice Hall.

Ghysels, E. and D. R. Osborn (2001). *The Econometric Analysis of Seasonal Time Series.* Cambridge University Press.

Gómez, V. and A. Maravall (2001). Automatic Modeling Methods for Univariate Series. In D. Peña, G. C. Tiao, and R. S. Tsay (Eds.), *A Course in Time Series Analysis*, pp. 171–201. New York: Wiley.

Harvey, A., S. J. Koopman, and M. Riani (1997). The Modeling and Seasonal Adjustment of Weekly Observations. *Journal of Business and Economic Statistics 15*(3), 354–368.

Holt, C. C. (2004). Forecasting Seasonals and Trends by Exponentially Weighted Moving Averages. *International Journal of Forecasting 20*(1), 5 – 10.

Hyndman, R. J. and Y. Khandakar (2008, July). Automatic Time Series Forecasting: The forecast Package for R. *Journal of Statistical Software 27*(3), 1–22.

Koopman, S. J., A. C. Harvey, J. A. Doornik, and N. Shephard (2009). *Structural Time Series Analyzer, Modeller, and Predictor (STAMP 8.2).* Timberlake Consultants Ltd.

Koopman, S. J. and M. Ooms (2003). Time Series Modelling of Daily Tax Revenues. *Statistica Neerlandica 57*(4), 439–469.

Koopman, S. J. and M. Ooms (2006). Forecasting Daily Time Series Using Periodic Unobserved Components Time Series Models. *Computational Statistics & Data Analysis 51*(2), 885–903.

Ladiray, D. and B. Quenneville (2001). *Seasonal Adjustment with the X-11 Method*, Volume 158. Springer Science & Business Media.

López, J. (2015). tsoutliers: R Package for Automatic Detection of Outliers in Time Series. Technical report.

Pierce, D. A., M. R. Grupe, and W. P. Cleveland (1984). Seasonal Adjustment of the Weekly Monetary Aggregates: A model-based Approach. *Journal of Business and Economic Statistics 2*(3), 260–270.

Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery (1992). *Numerical Recipes in C: The Art of Scientific Computing.* Cambridge university press.

Tukey, J. W. (1960). A Survey of Sampling from Contaminated Distributions. In I. Olkin (Ed.), *Contributions to probability and statistics : essays in honor of Harold Hotelling*, pp. 448–485. Stanford, California: Stanford University Press.

Venables, W. N. and B. D. Ripley (2002). *Modern Applied Statistics with S* (4 ed.). Springer.

Verbesselt, J., R. Hyndman, G. Newnham, and D. Culvenor (2010). Detecting Trend and Seasonal Changes in Satellite Image Time Series. *Remote sensing of Environment 114*(1), 106–115.

Vrieze, S. I. (2012). Model Selection and Psychological Theory: a Discussion of the Differences between the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC). *Psychological methods 17*(2), 228–243.

Young, M. R. (1996). Robust Seasonal Adjustment by Bayesian Modelling. *Journal of Forecasting 15*(5), 355–367.

A Example of Seasonal Adjustment with BAYSEA

Results for Seasonal Adjustment based on the BAYSEA Algorithm
for Simulated Time Series

Figure 5

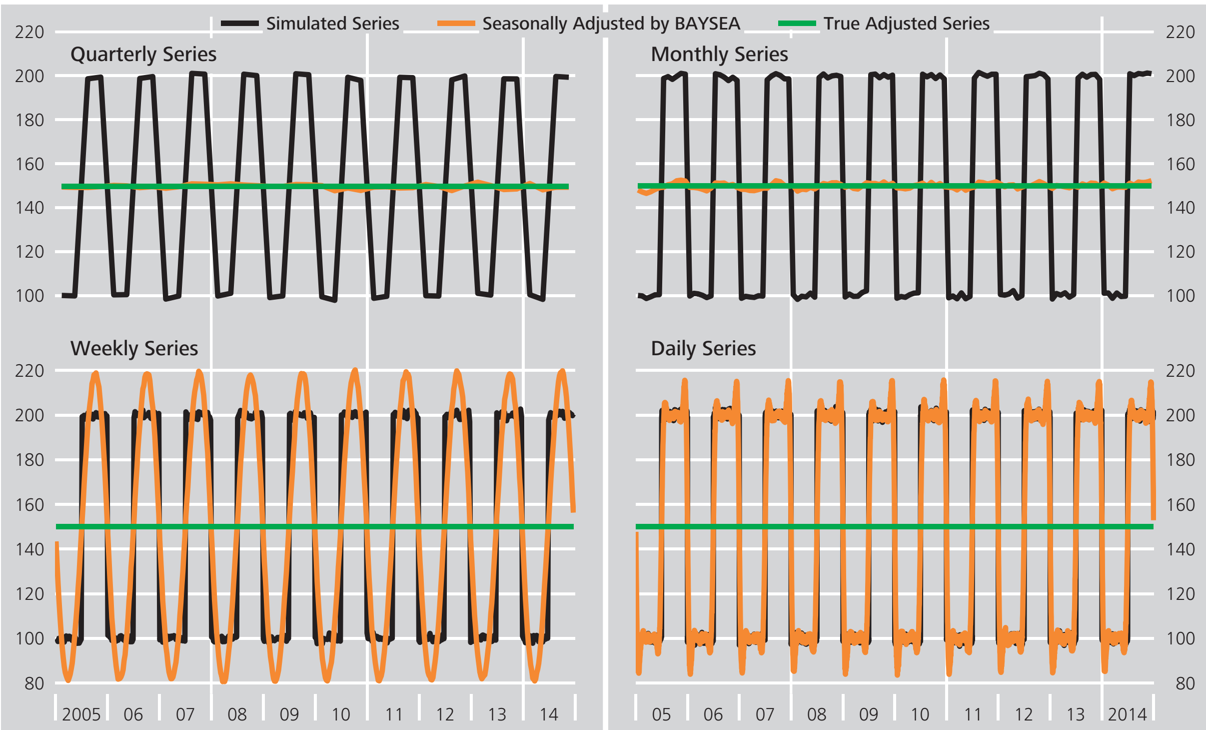

Deutsche Bundesbank

The so-called bayesian seasonal adjustment, BAYSEA, was developed as a structural model based alternative to existing methods such as Census X-11 for the seasonal adjustment of quarterly and monthly time series (Akaike, 1980; Akaike and Ishiguro, 1983). Again, the idea is to decompose a time series Y_t with seasonal frequency f into a trend-cycle T_t, a seasonal S_t and an irregular component I_t (see Equation 1) based on assumptions of the behaviour of the components. The systematic part, i.e. S_t and T_t, should not deviate systematically from the time series Y_t, so that $(Y_t - T_t - S_t)^2$ ought to be small. The assumption of a locally smooth behaviour of the trend translates into the minimisation of the second difference of T_t. Finally, the seasonal component is supposed to be stable. Simultaneously, this can be achieved by minimising

$$\begin{aligned}((Y_t - T_t - S_t)^2 + \kappa_1^2\{\kappa_2^2[(T_t - T_{t-1}) - (T_{t-1} - T_{t-2})]^2 +\\ (S_t - S_{t-f})^2 + \kappa_3^2(S_t + S_{t-1} + \ldots + S_{t-f+1})^2\})\end{aligned}$$

using constrained least squares, where κ_1, κ_2, and κ_3 are constants, usually determined by using information criteria.

As Figure 5 shows, the main drawback of BAYSEA is its need for very long time series, when the number of observations per year increases. Therefore it does not lend itself easily to the seasonal adjustment of daily time series.

B RegARIMA Results

Table 3: RegARIMA Results for Currency in Circulation in Germany. Based on Time Series only Adjusted for Intra-Weekly Seasonal Effects.

Regressor	Coefficient	S.E.	T-Value	Regressor	Coefficient	S.E.	T-Value
$ma1$	0.08910	0.03018	2.95	S_{22}	0.00011	0.00006	1.93
$drift$	0.00023	0.00003	8.26	C_{22}	−0.00007	0.00006	−1.28
S_1	−0.00264	0.00121	−2.18	S_{23}	0.00008	0.00005	1.41
C_1	0.00325	0.00129	2.51	C_{23}	−0.00016	0.00005	−3.02
S_2	0.00252	0.00060	4.18	S_{24}	0.00020	0.00005	3.84
C_2	0.00124	0.00061	2.04	C_{24}	0.00028	0.00005	5.58
S_3	0.00230	0.00040	5.75	S_{25}	−0.00001	0.00005	−0.29
C_3	0.00087	0.00040	2.17	C_{25}	−0.00004	0.00005	−0.90
S_4	0.00159	0.00030	5.27	S_{26}	−0.00010	0.00005	−2.11
C_4	−0.00060	0.00030	−1.98	C_{26}	−0.00004	0.00005	−0.76
S_5	0.00164	0.00024	6.80	S_{27}	−0.00012	0.00005	−2.51
C_5	−0.00112	0.00024	−4.63	C_{27}	−0.00001	0.00005	−0.16
S_6	0.00086	0.00020	4.29	S_{28}	−0.00008	0.00005	−1.66
C_6	−0.00093	0.00020	−4.61	C_{28}	−0.00004	0.00004	−0.78
S_7	0.00015	0.00017	0.87	Eas{-7}	0.00033	0.00024	1.39
C_7	−0.00117	0.00017	−6.77	Eas{-6}	0.00041	0.00035	1.19
S_8	−0.00033	0.00015	−2.22	Eas{-5}	0.00046	0.00042	1.10
C_8	−0.00096	0.00015	−6.37	Eas{-4}	0.00125	0.00047	2.68
S_9	−0.00046	0.00013	−3.41	Eas{-3}	0.00225	0.00050	4.50
C_9	−0.00054	0.00013	−4.01	Eas{-2}	0.00347	0.00052	6.69
S_{10}	−0.00040	0.00012	−3.34	Eas{-1}	0.00486	0.00053	9.25
C_{10}	−0.00007	0.00012	−0.57	Eas{0}	0.00437	0.00052	8.42
S_{11}	−0.00043	0.00011	−3.92	Eas{1}	0.00436	0.00050	8.73
C_{11}	0.00011	0.00011	1.05	Eas{2}	0.00436	0.00047	9.36
S_{12}	0.00091	0.00010	9.03	Eas{3}	0.00497	0.00042	11.94
C_{12}	0.00055	0.00010	5.48	Eas{4}	0.00331	0.00035	9.56
S_{13}	−0.00013	0.00009	−1.41	Eas{5}	0.00079	0.00024	3.27
C_{13}	0.00041	0.00009	4.38	Asc{-7}	−0.00072	0.00024	−3.02
S_{14}	−0.00003	0.00009	−0.40	Asc{-6}	−0.00141	0.00035	−4.08
C_{14}	0.00040	0.00009	4.63	Asc{-5}	−0.00120	0.00042	−2.88
S_{15}	0.00012	0.00008	1.44	Asc{-4}	−0.00062	0.00047	−1.33
C_{15}	0.00025	0.00008	3.07	Asc{-3}	0.00005	0.00050	0.09
S_{16}	0.00028	0.00008	3.71	Asc{-2}	0.00031	0.00052	0.61
C_{16}	0.00002	0.00008	0.26	Asc{-1}	0.00057	0.00053	1.09
S_{17}	0.00033	0.00007	4.60	Asc{0}	0.00078	0.00052	1.49
C_{17}	−0.00004	0.00007	−0.60	Asc{1}	0.00223	0.00050	4.47
S_{18}	0.00024	0.00007	3.51	Asc{2}	0.00129	0.00047	2.77
C_{18}	−0.00016	0.00007	−2.32	Asc{3}	−0.00002	0.00042	−0.05
S_{19}	0.00012	0.00007	1.77	Asc{4}	−0.00028	0.00035	−0.80
C_{19}	−0.00012	0.00006	−1.85	Asc{5}	−0.00012	0.00024	−0.51
S_{20}	0.00011	0.00006	1.74	LS12.12.02	0.00947	0.00025	37.65
C_{20}	−0.00013	0.00006	−2.12	TC12.12.03	−0.00401	0.00022	−18.14
S_{21}	−0.00002	0.00006	−0.27	AO13.01.03	0.00122	0.00015	8.23
C_{21}	−0.00014	0.00006	−2.31				

Model order: (0 1 1), trigonometric sinusoid (S_j) and cosinusoid (C_j) regressors with j cycles per year, Eas{m} is the effect m days after Easter, Asc{n} is effect n days after Ascension. ZZyy.mm.dd is an outlier of type ZZ on the following day: year yy, month mm and day dd.

C Comparison of Seasonal Adjustment Procedures

Comparison of Seasonal Adjustment Procedures based on End-of-Month Figures

Figure 6

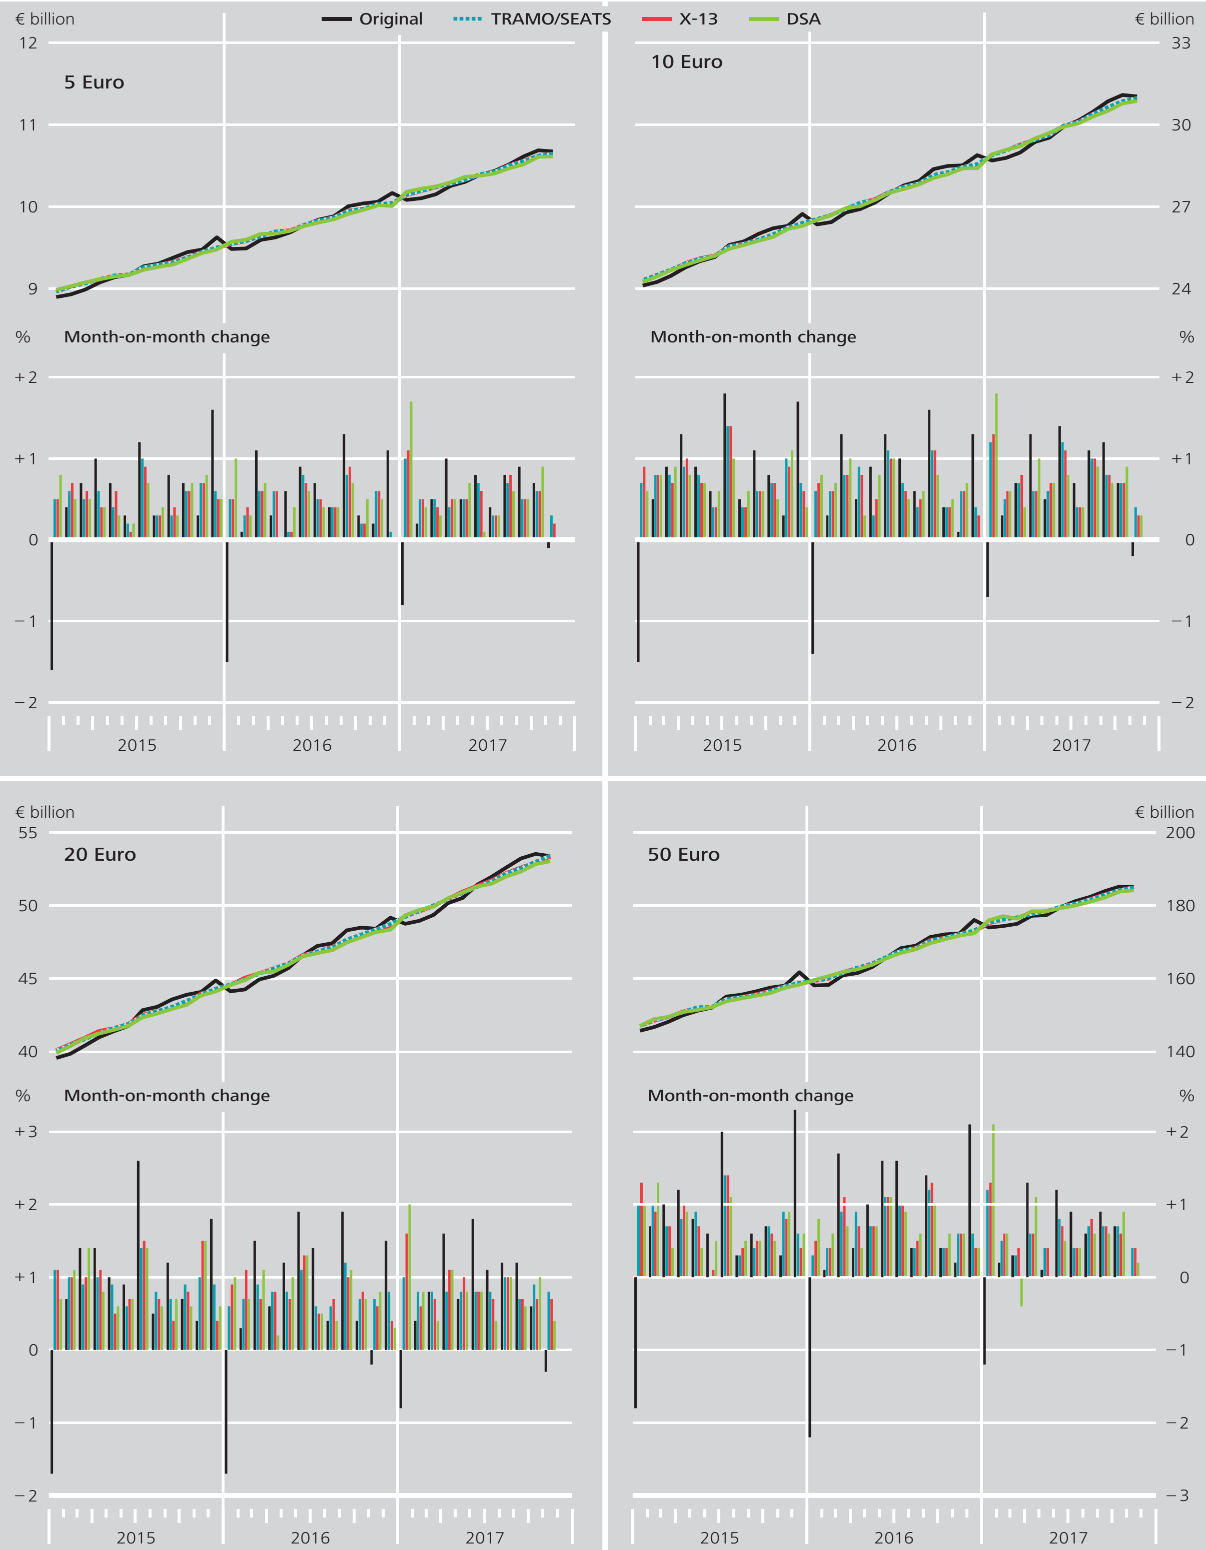

Deutsche Bundesbank

Comparison of Seasonal Adjustment Procedures based on End-of-Month Figures (cont.)

Figure 6

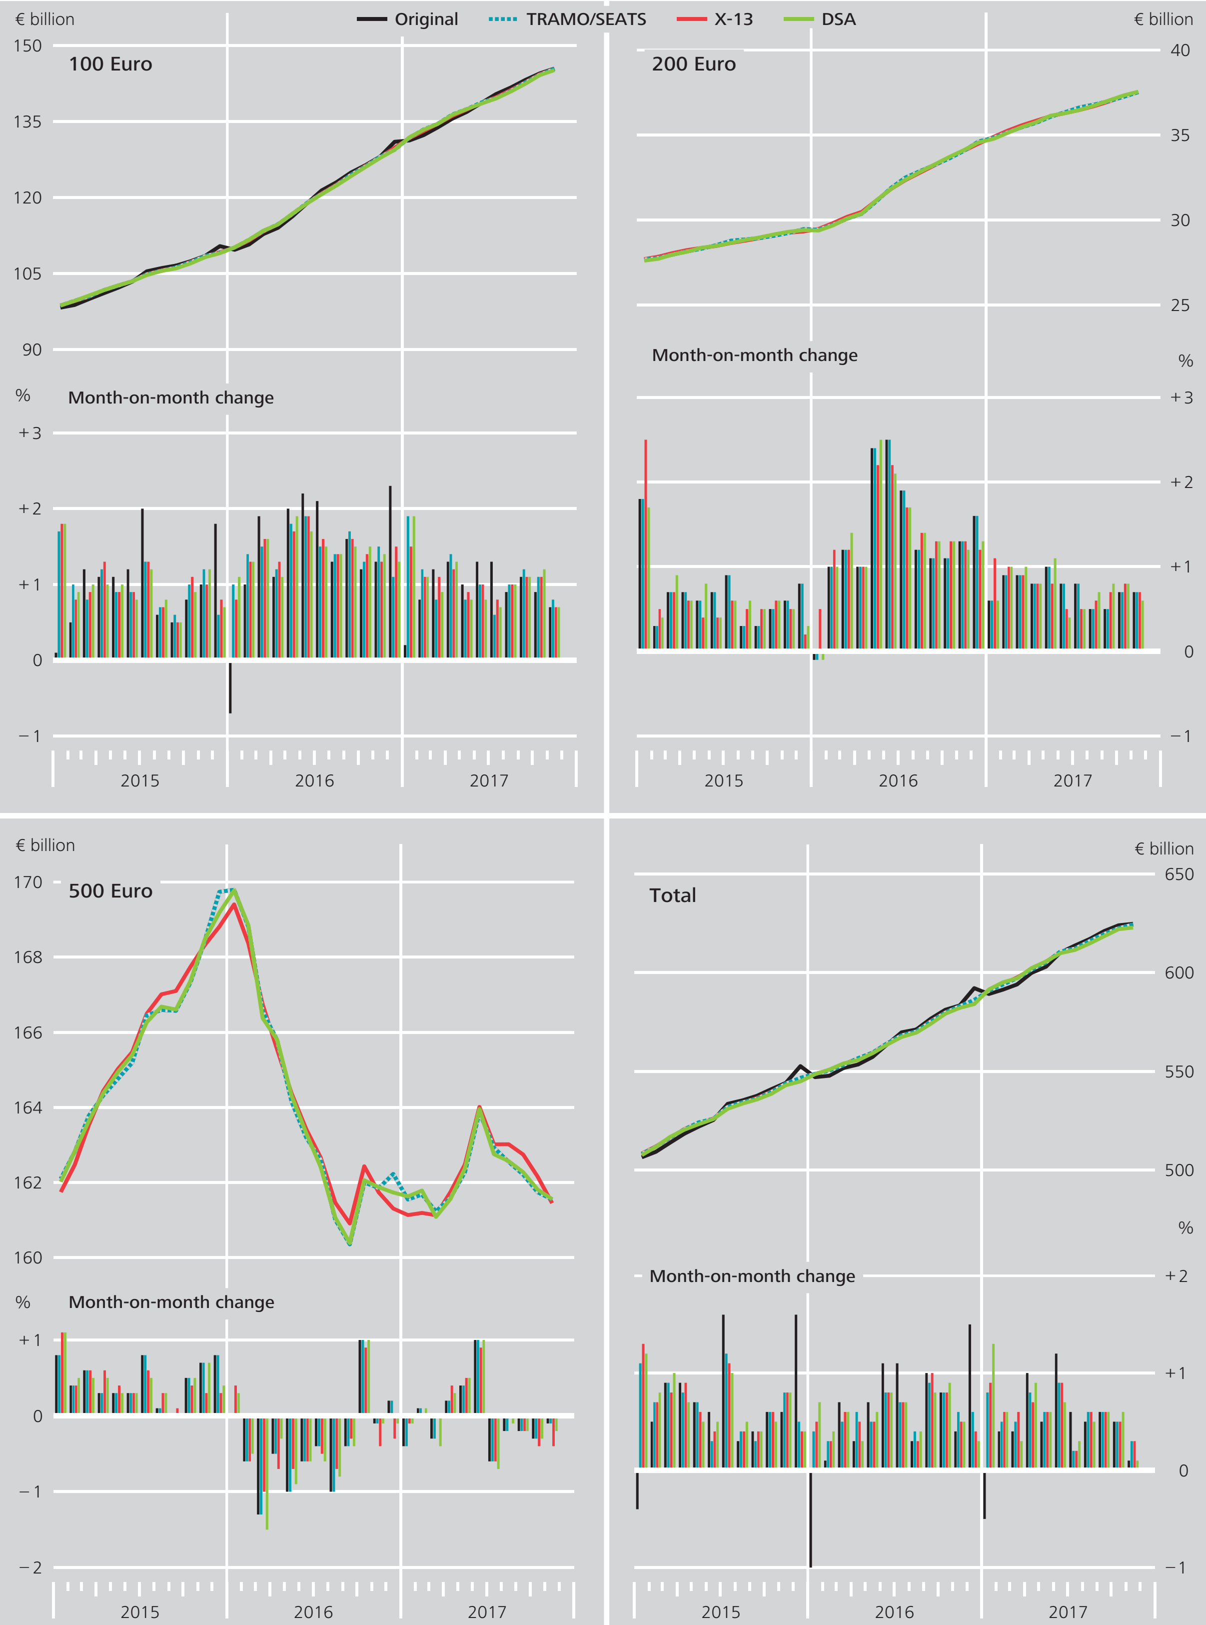

Deutsche Bundesbank

Comparison of Seasonal Adjustment Procedures based on Monthly Means

Figure 7

Deutsche Bundesbank

Comparison of Seasonal Adjustment Procedures based on Monthly Means (cont.)

Figure 7

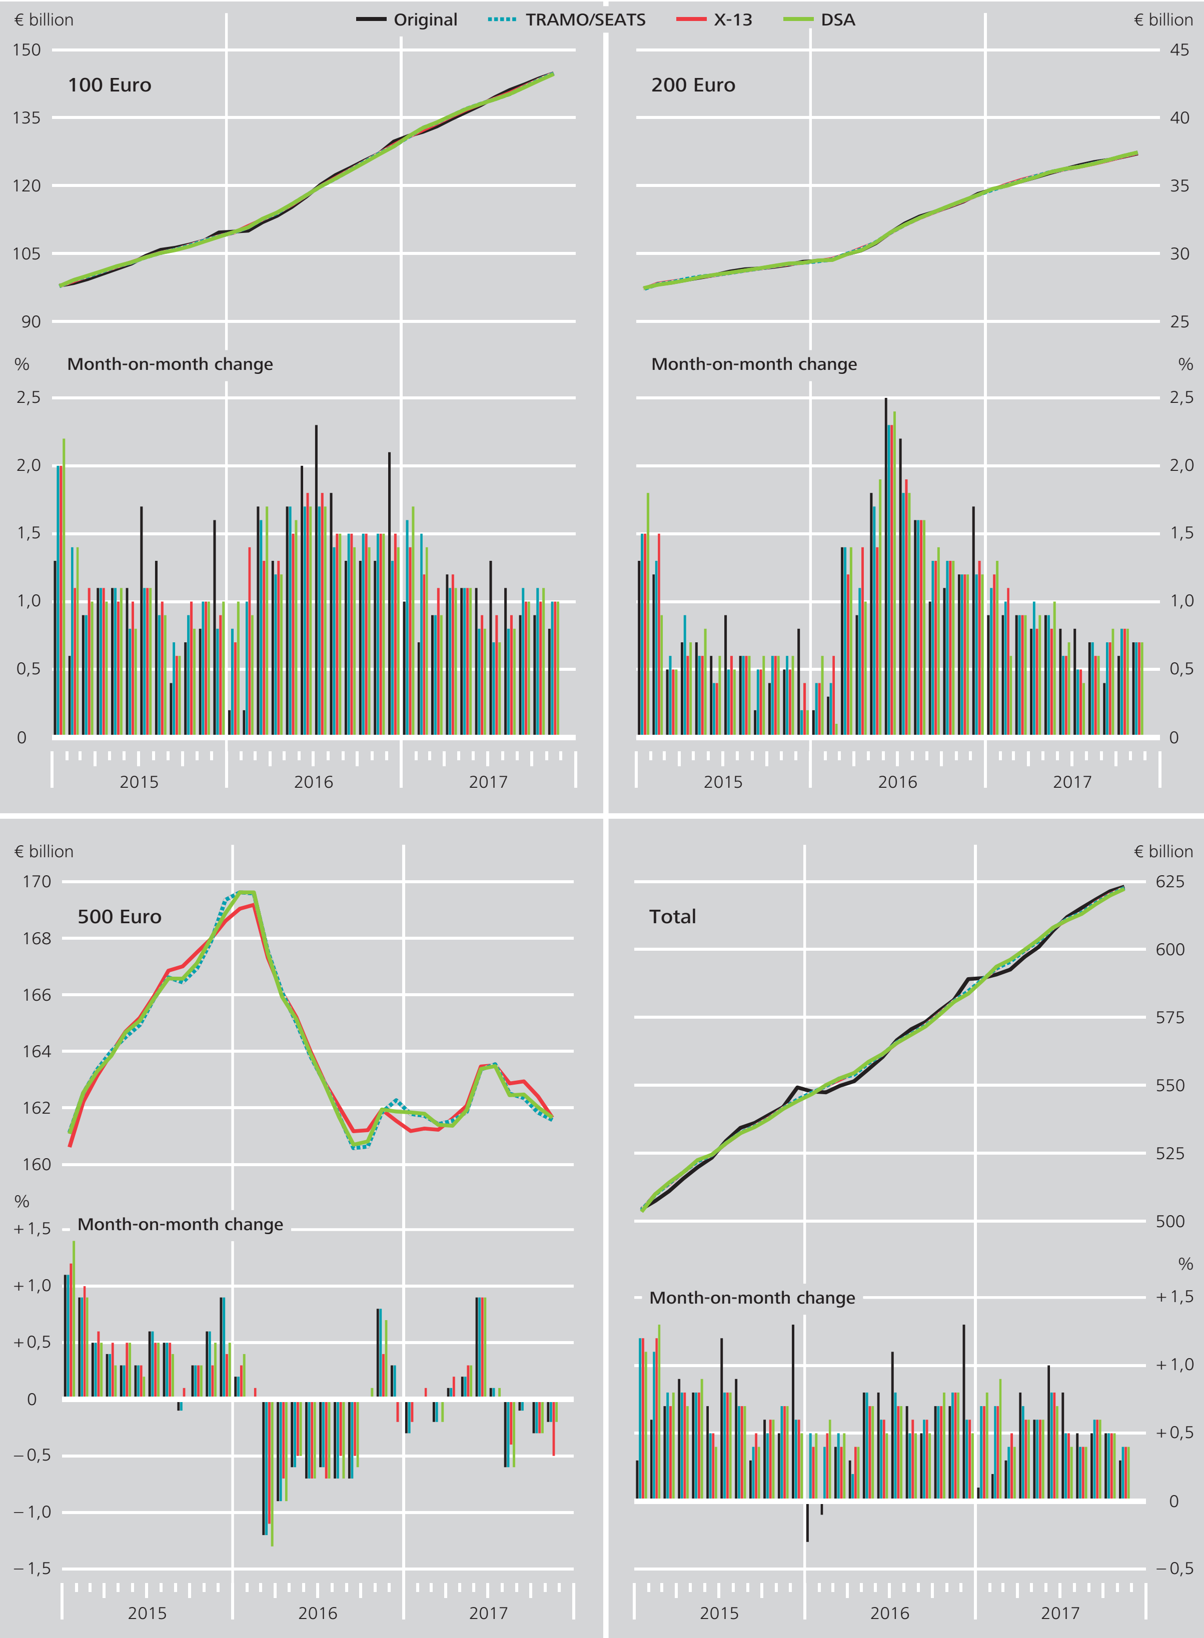

Deutsche Bundesbank

Comparison of Seasonal Adjustment Procedures based on End-of-Month Figures of Simulated Time Series

Figure 8

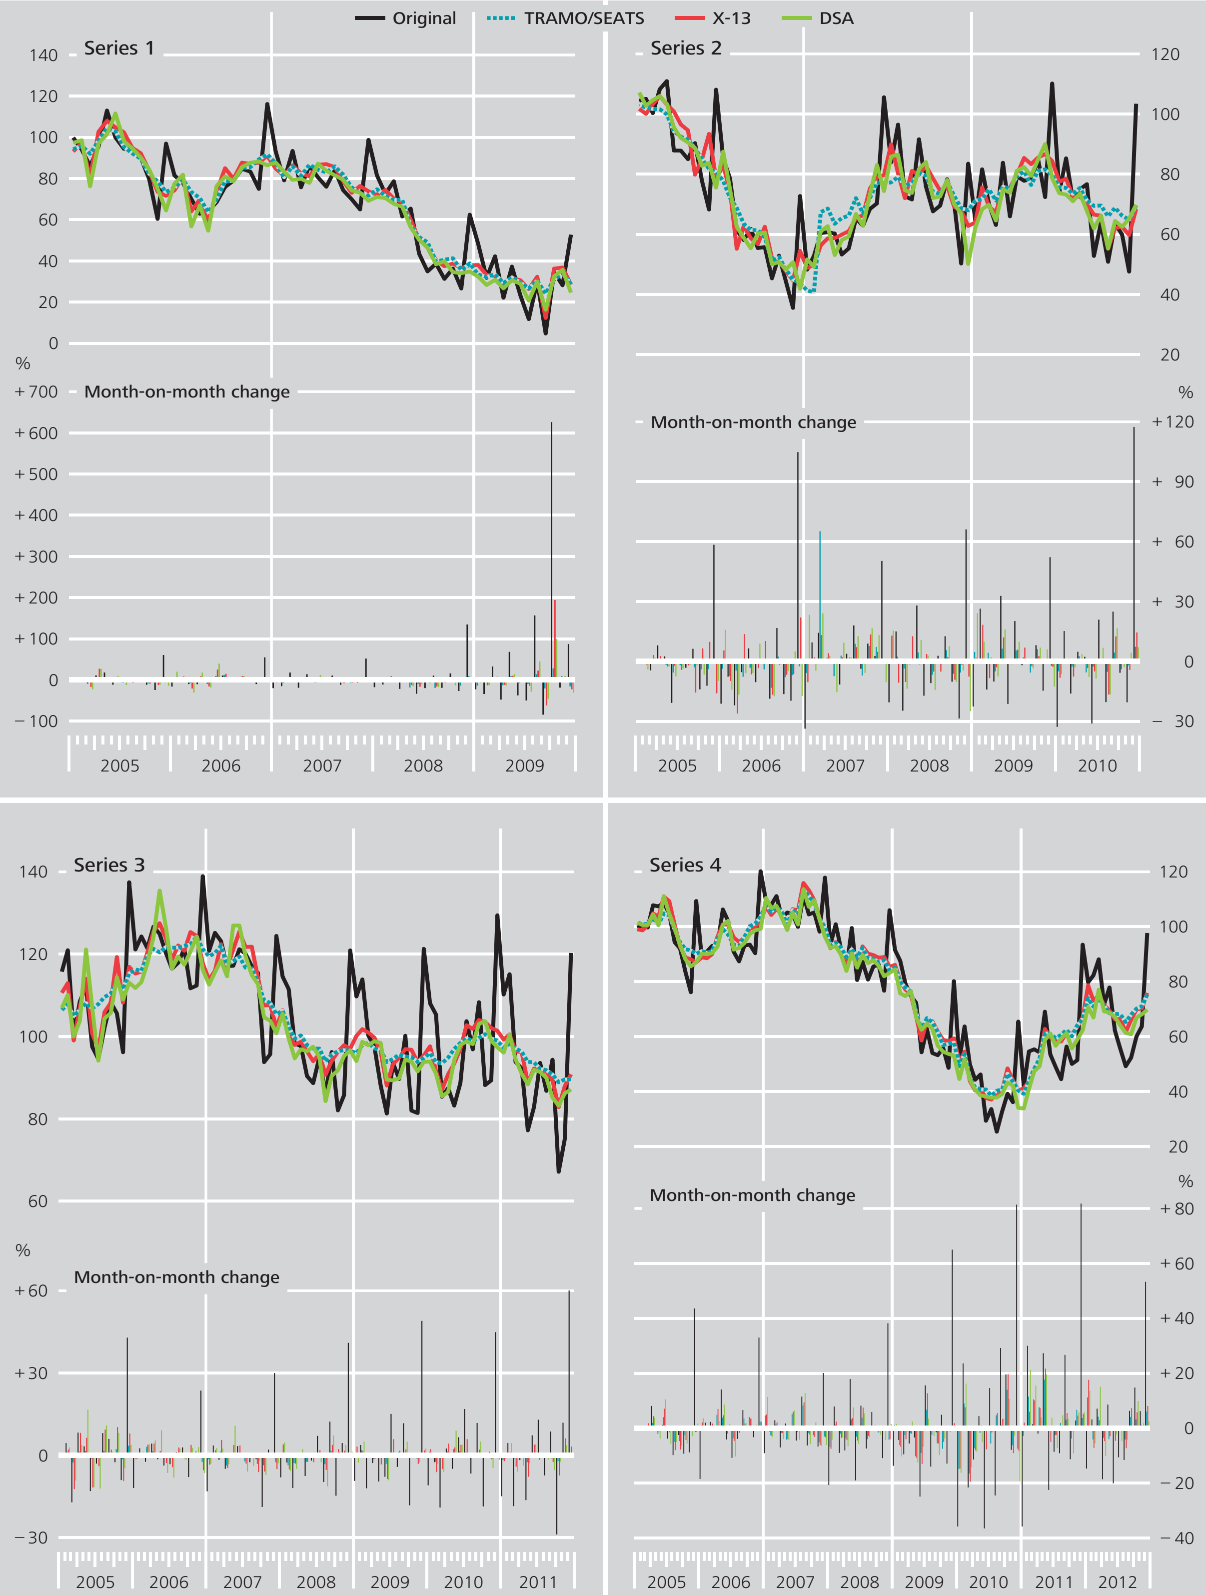

Deutsche Bundesbank

Comparison of Seasonal Adjustment Procedures based on End-of-Month Figures of Simulated Time Series (cont.)

Figure 8

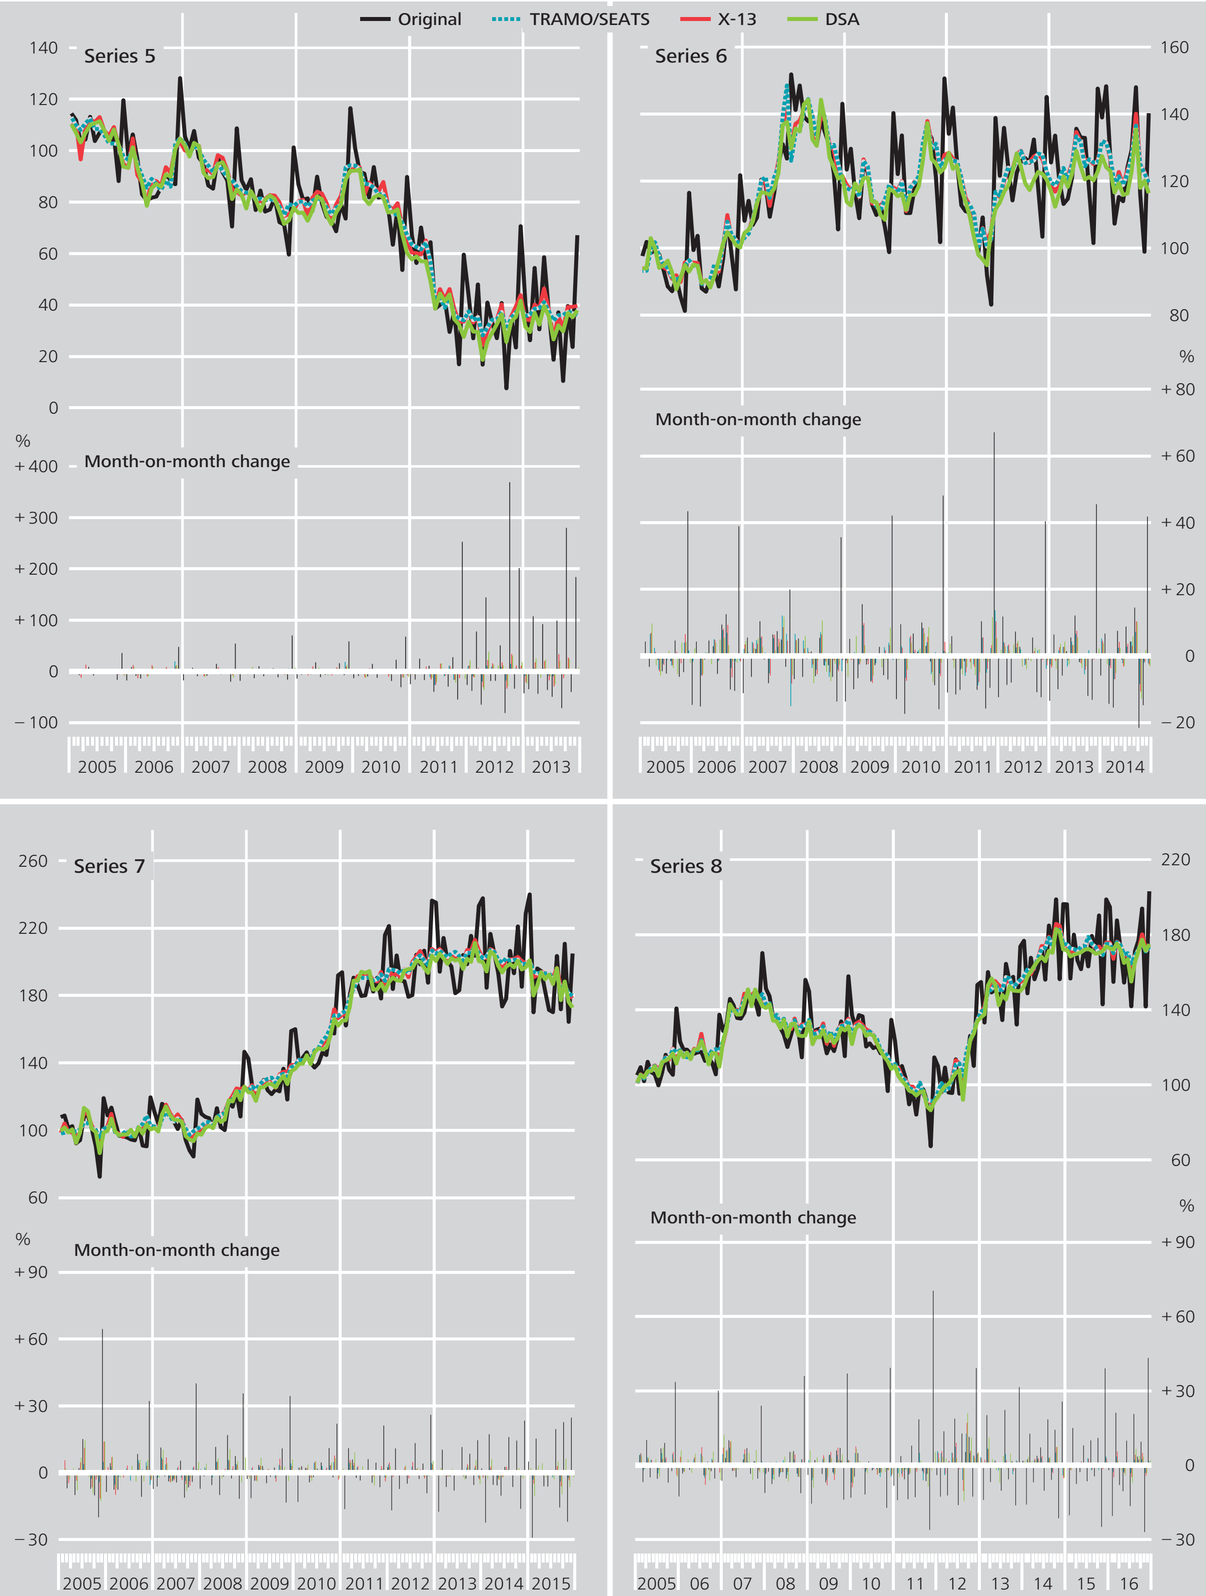

Deutsche Bundesbank

Comparison of Seasonal Adjustment Procedures based on Monthly Means of Simulated Time Series

Figure 9

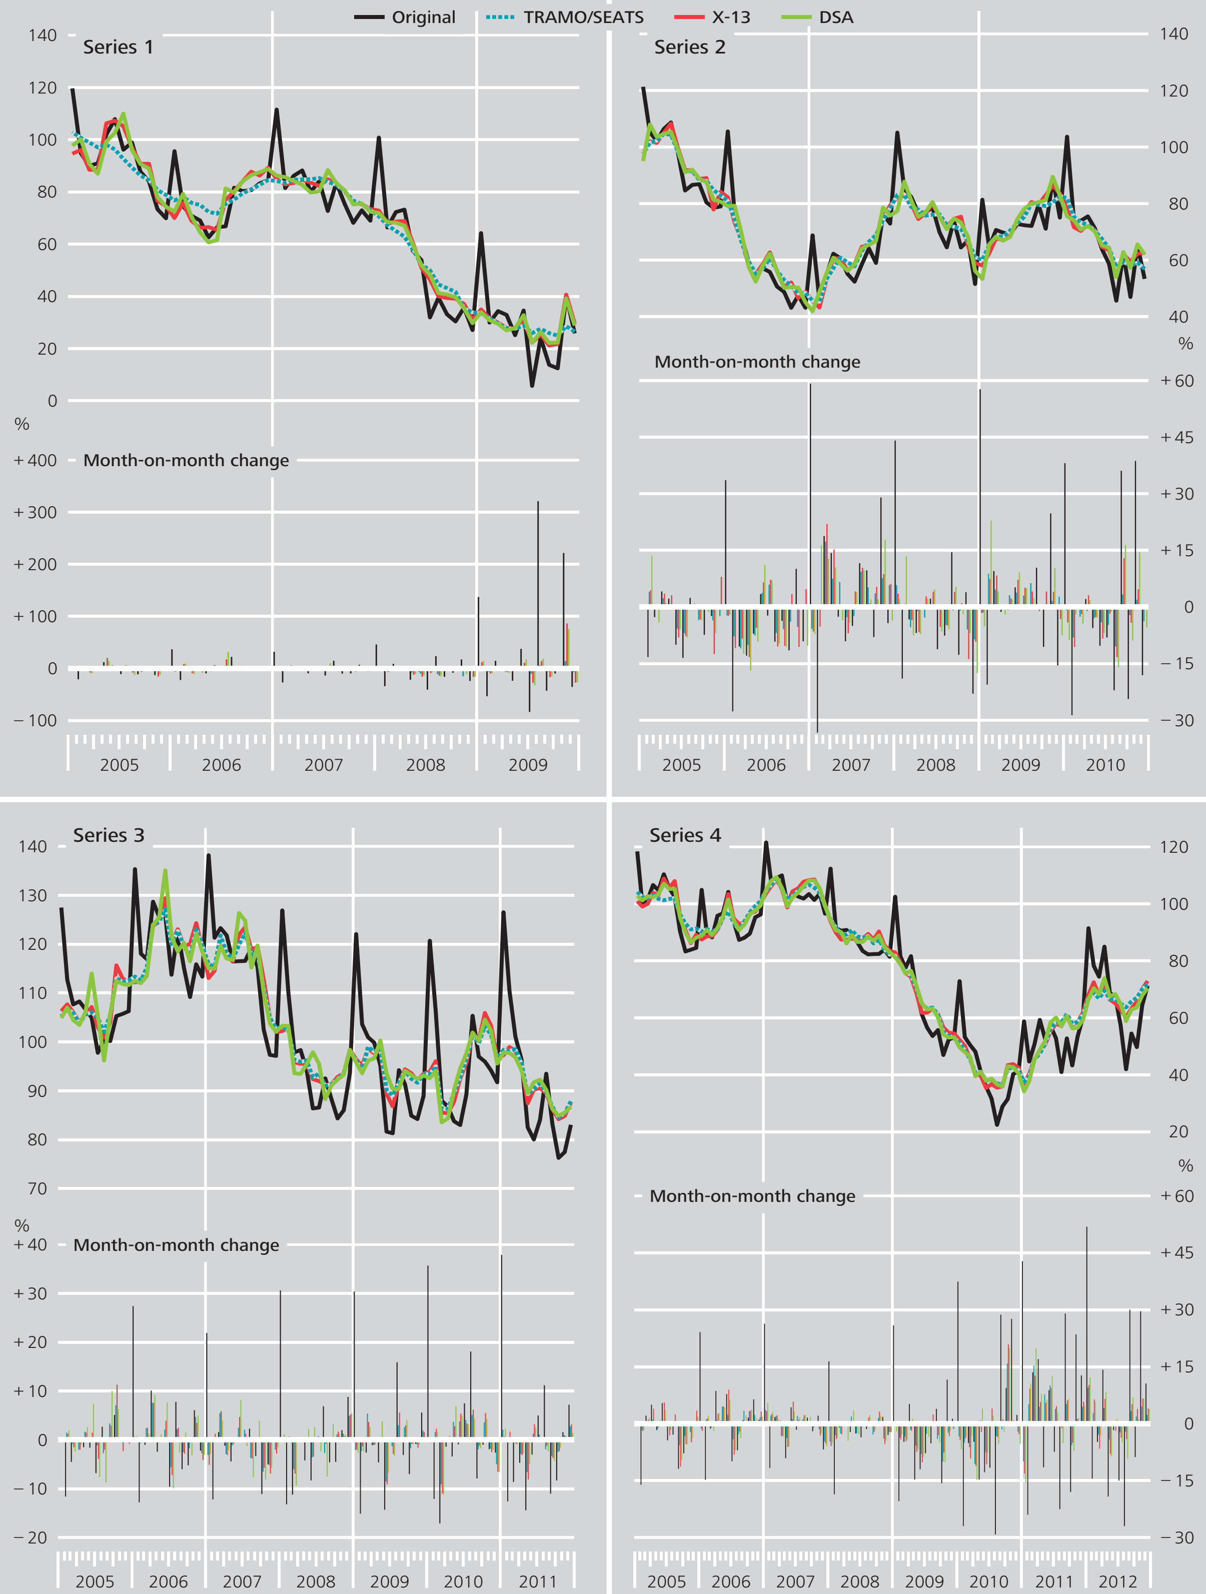

Deutsche Bundesbank

Comparison of Seasonal Adjustment Procedures based on
Monthly Means of Simulated Time Series (cont.)

Figure 9

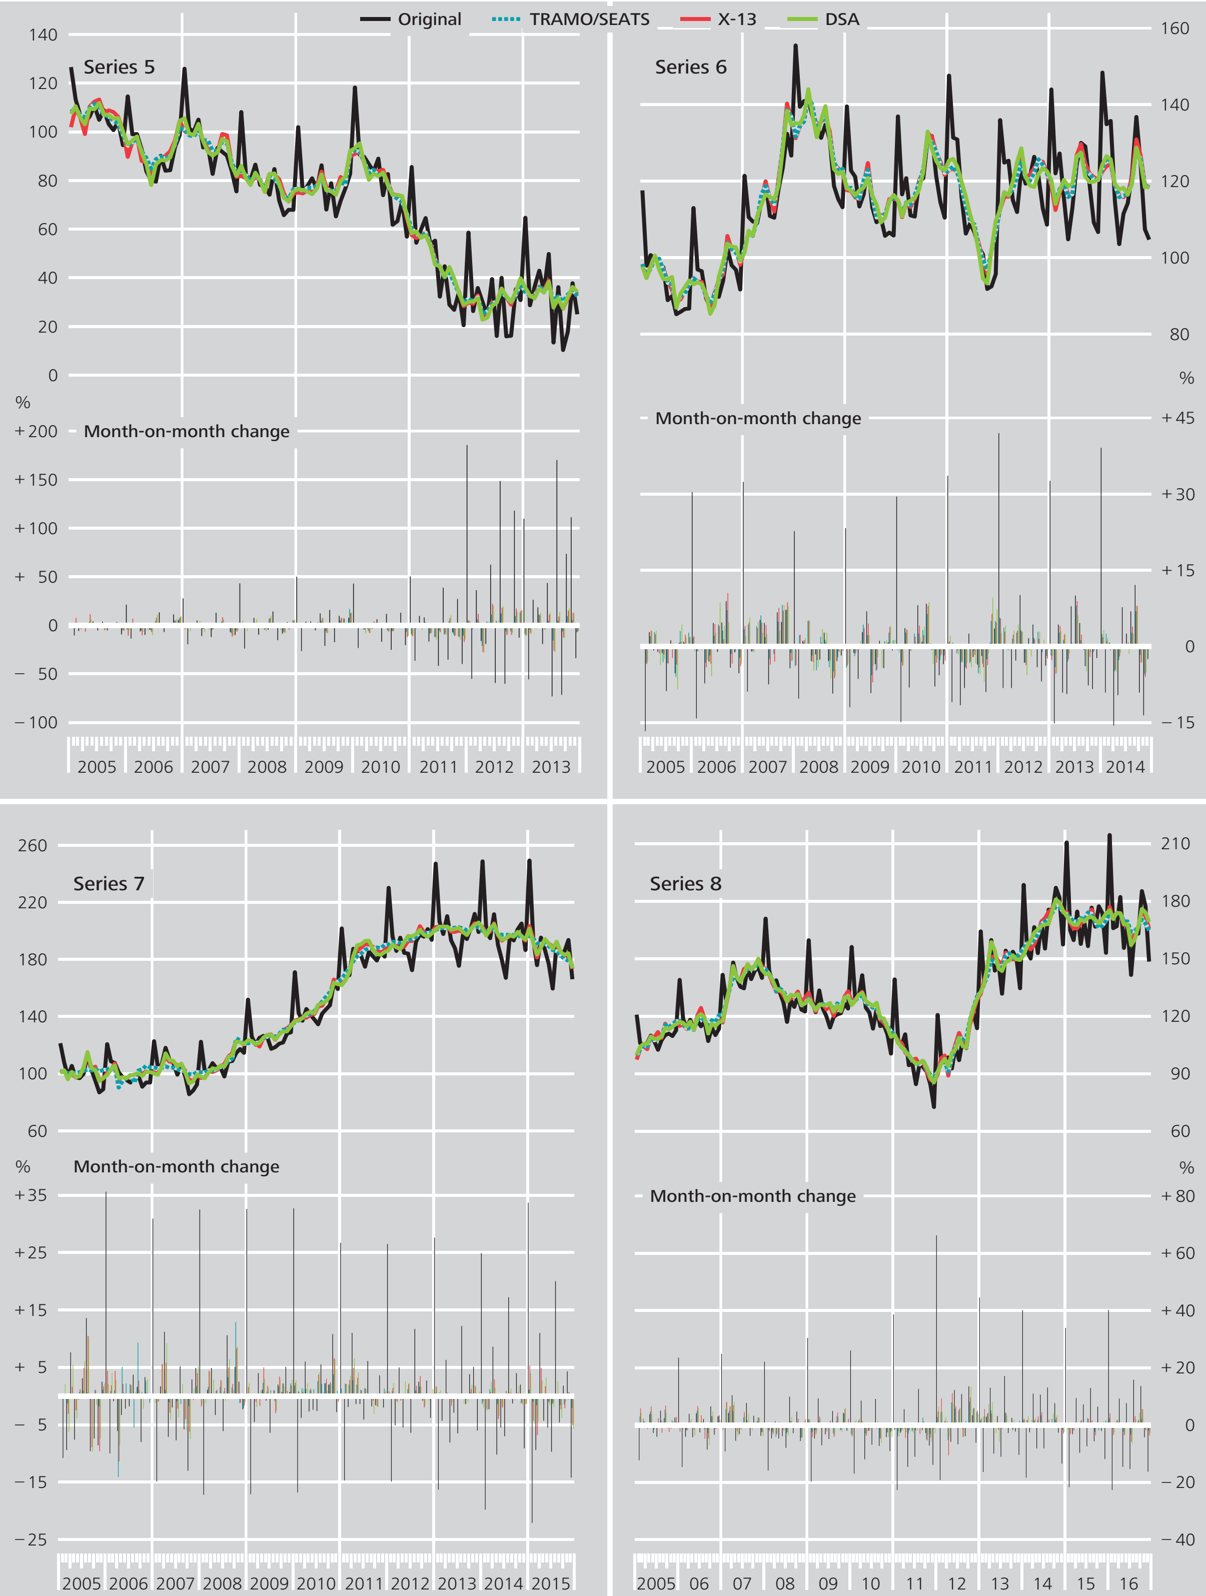

Deutsche Bundesbank

Comparison of Seasonal Adjustment Procedures based on
Daily Figures of Simulated Time Series

Figure 10

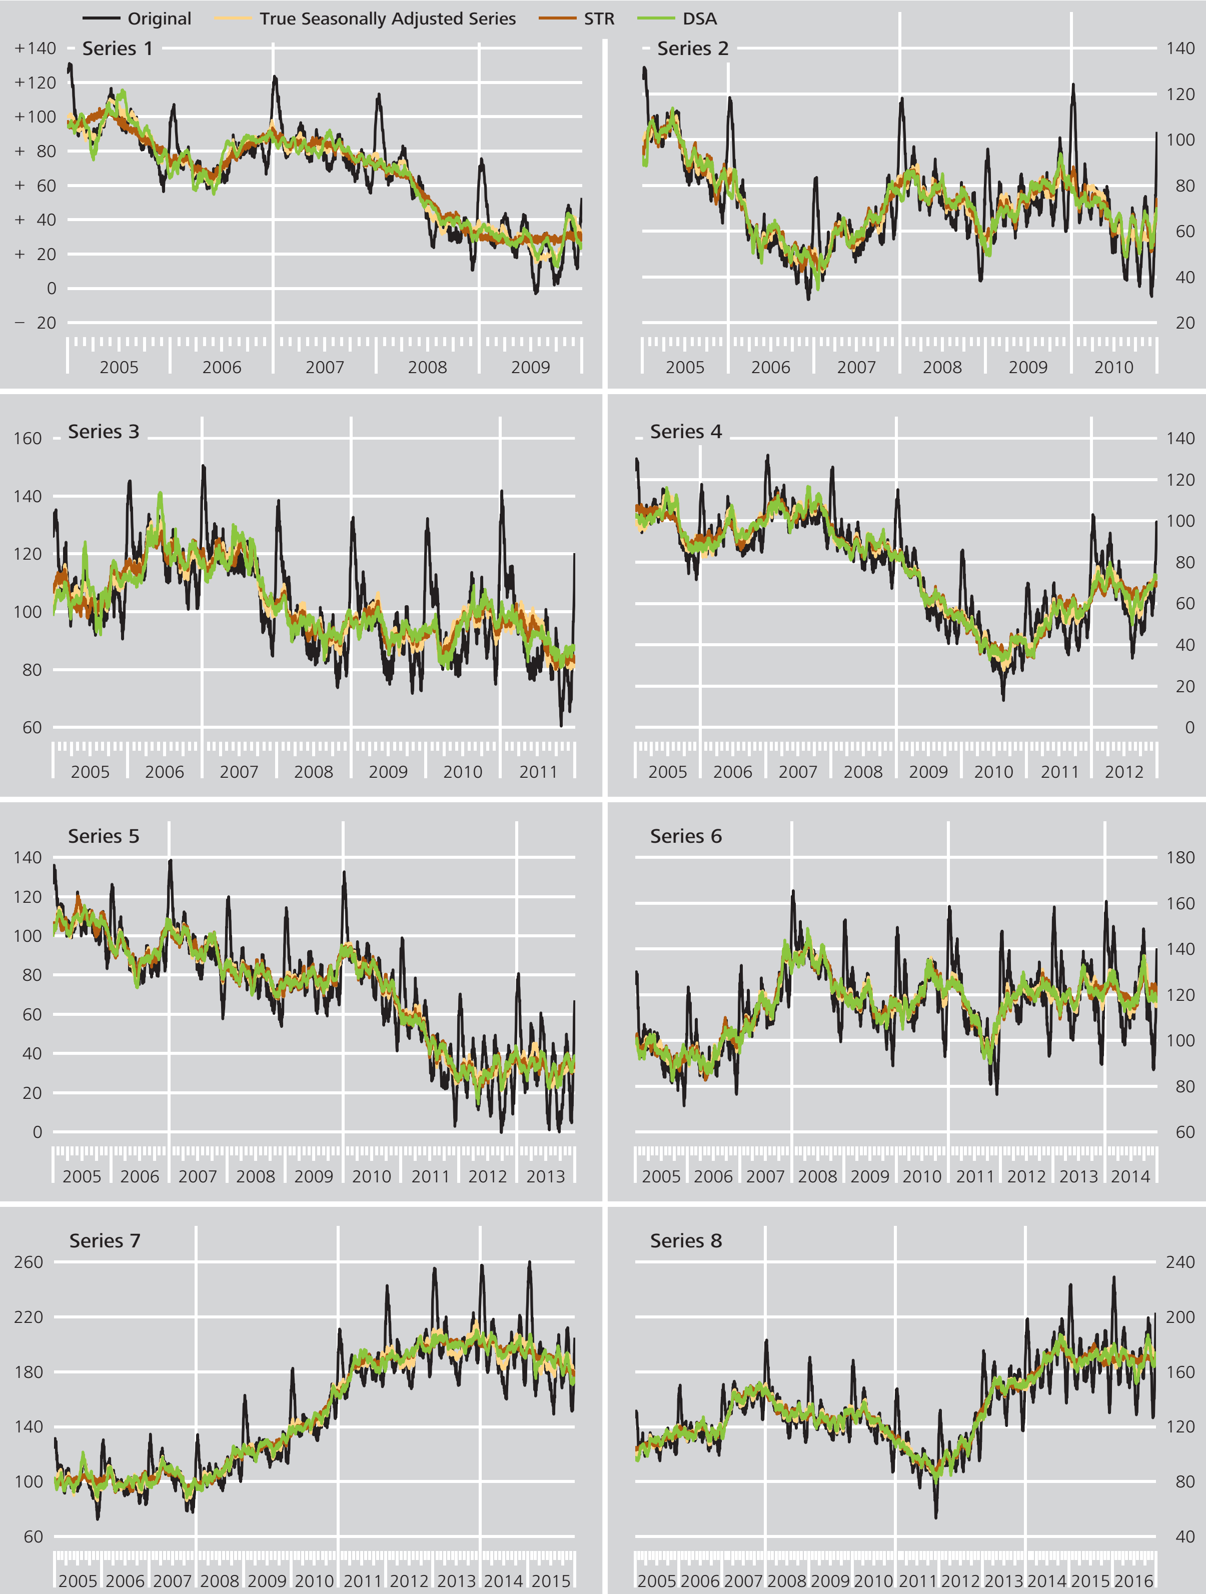

Deutsche Bundesbank