Bank of Russia

Seasonal adjustment of the Bank of Russia Payment System financial flows data

WORKING PAPER SERIES

No. 65 / December 2020

Sergey Seleznev, Natalia Turdyeva,
Ramis Khabibullin, AnnaTsvetkova

Sergey Seleznev
Bank of Russia. Email: SeleznevSM@mail.cbr.ru.

Natalia Turdyeva
Bank of Russia. Email: TurdyevaNA@mail.cbr.ru.

Ramis Khabibullin
Bank of Russia. Email: KhabibullinRA@mail.cbr.ru.

Anna Tsvetkova
Bank of Russia. Email: TsvetkovaAN@mail.cbr.ru

Bank of Russia Working Paper Series is anonymously refereed by members of the Bank of Russia Research Advisory Board and external reviewers.

Cover image: Shutterstock.com

Address: 12 Neglinnaya street, Moscow, 107016
Tel.: +7 495 771-91-00, +7 495 621-64-65 (fax)
Website: www.cbr.ru

Abstract

This paper describes the seasonal adjustment algorithm used by the Bank of Russia to clean up data for ‘Monitoring of Sectoral Financial Flows’ weekly publication. We have developed a simple and fast procedure based on a set of trigonometric functions and dummy variables that demonstrates good results in terms of various quality metrics and can be easily modified for working with more flexible model specifications.

JEL classification: C11, C22, E32, E37.

Keywords: daily seasonal adjustment, time series, sectoral financial flows, Bayesian estimator.

Contents

1. Introduction

Due to the deterioration of economic conditions caused by the spread of the coronavirus infection, economists and statisticians across the globe have faced the need to monitor the state of the economy at weekly, and even daily, frequency. The unprecedented pace of decline in economic activity demanded that governments of all countries provide energetic and informed responses.

Unconventional high-frequency data have already attracted the attention of economists working in the field of economic policy for a number of years (e.g. Hueng ed. (2020)). Based on these pre-existing developments, many studies aimed at using large amounts of high-frequency data to analyze the emerging economic conditions in the situation of the economic crisis caused by the pandemic, (e.g. Carvalho et al. (2020), Chetty et al. (2020), Lewis et al. (2020)).

Over this period, the Bank of Russia developed a set of indicators based on Bank of Russia Payment System (BRPS) data, which made it possible to observe the dynamics of sectoral financial flows nearly in real-time. The analysis of these indicators posed a number of challenges; one of them was removing seasonal effects, as they greatly complicate understanding the reasons behind changes in indicator dynamics (Figure 1 in Appendix A). This paper describes the methodology for isolating the seasonal component in the financial flows data.

The problem of seasonal component extraction is a special case of the problem of decomposing a time series into various components, and is well studied in the literature. However, it is necessary to recognise that in its simplest formulation, this problem is not well defined. While seemingly obvious, the decomposition of time series will depend on the researcher's definitions of trend, cycle, and seasonality, and on the way they are modelled. Unfortunately, there is no objective criterion for measuring the quality of the decomposition (what one person considers to be a correct decomposition, another person may not). In most cases, seasonal adjustment is applied to solve a specific problem. Hence, the methodology should be appropriate for solving it. With these considerations in mind, we chose the method of seasonal adjustment. In our case, the aim of decomposition was to clean up financial flows data of a set of strictly periodic patterns that are uninformative for further analysis. Essentially, this set of patterns is the definition of seasonality used in this paper, and the degree of its removal serves as the final point of reference for choosing the seasonal adjustment procedure.

In practice, we encountered a number of additional limitations that also affected the choice of procedure currently applied to adjusting financial flows data. They included the type of patterns excluded from the raw data, the specifics of financial flows data, as well as the speed of algorithms and the ability to adjust the set of excluded patterns quickly. Due to the daily periodicity of data and our specific definition of seasonality, we were not able to use the algorithms widely applied in economic analysis. At the same time, lack of transactions on weekends and holidays made it difficult to use well-known packages such as Prophet (Taylor and Letham (2017)). Using the ideas of Taylor and Letham (2017) and the specifics of financial flows data, and taking into account the limitations, such as the speed and the ability to adjust the set of excluded patterns quickly, we have developed a basic seasonal adjustment procedure, which is currently used to clean up data for the weekly report published on the Bank of Russia's website.[1]

The basic procedure, potential extensions, and a discussion of methodology are presented in Section 2. Section 3 discusses criteria that help assess the quality of seasonal adjustment. Section 4 describes the sectoral financial flows data. Section 5 provides the results. Section 6 discusses relevant literature. Section 7 concludes.

2. Seasonal adjustment

Basic procedure

As part of the basic procedure, we identify a set of multiplicative seasonal patterns (s_t), which, similarly to seasonal patterns identified in the Facebook Prophet procedure, consists of an intraweekly (s_t^w) and intra-annual (s_t^y) components. In order to take the intramonthly seasonality present in the data into account explicitly (Figure 1), we also add the intramonthly component (s_t^m). Thus, the seasonal component is modelled as:[2]

$$s_t = s_t^w + s_t^m + s_t^y \quad (1)$$

We assume that the remaining part is the sum of the non-stationary (tr_t) and stationary components (e_t) that do not contain seasonality[3]; they will be referred to as trend and residuals. Taking into account that there are days with no payments (weekends and holidays), our model can be presented as follows:

$$Y_t = \begin{cases} e^{s_t^w + s_t^m + s_t^y + tr_t + e_t} \ if \ there \ are \ payments \\ 0 \ otherwise \end{cases} \quad (2)$$

where Y_t is the data after preliminary processing. Hereinafter, for simplicity, we will use x_t to denote the logarithm of the variable X_t and concentrate only on the non-zero part without loss of generality.

From the point of view of asymptotic theory, the allocation of a non-stationary component in type (2) models should play a key role in the sense that incorrect trend specifications almost always lead to an inadequate estimation, while errors in the specification of residuals often do not cause asymptotic problems. Given this fact, the

[1] 'Monitoring of Sectoral Financial Flows' is a weekly analytical publication of the Bank of Russia providing aggregate information on payments that passed through the BRPS (http://www.cbr.ru/analytics/finflows/).
[2] We do not take the effect of holidays into account since it does not significantly affect the estimation of the intramonthly, daily, and intra-annual components; however, it can easily be included in the model if necessary, as part of its extension.
[3] Formally, we assume that:

$$\frac{\sum_{t=1}^{T} S_t}{T} \to 0$$

$$\frac{\sum_{t=1}^{T} S_t E_t}{T} \to 0 \ a.s$$

$$\exists K: \forall k \geq K \ \frac{\sum_{t=1}^{T} \Delta^k S_t \Delta^k T_t}{T} \to 0 \ a.s$$

stationary component is modelled as a normal distribution with zero mean and estimated variance (σ_e^2):

$$e_t \sim N(0, \sigma_e^2) \qquad (3)$$

Preliminary visual analysis did not reveal any strict non-linearities in the pre-pandemic data. Since the ultimate goal of the model was not to make a forecast for the subsequent period, 2020 was excluded from the estimation. The trend component was modelled as a linear dependence on time:

$$tr_t = \theta^{tr} t \qquad (4)$$

where θ^{tr} is the estimated parameter.

The intraweekly component is estimated using dummy variables for each day of the week, which allows removing all kinds of periodicity on a daily basis:

$$s_t^w = \sum_{k=1}^{5} \theta_k^w I_t^k \qquad (5)$$

where I_t^k is an indicator that equals 1 for business day k; θ_k^w are estimated parameters.

Similar to many papers on seasonal adjustment (e.g. Taylor and Letham (2017)), periodically repeated patterns where the number of days in a period is large are removed not by using dummy variables but rather by using a set of trigonometric functions that can approximate any periodic function, given their adequately large number:

$$s_t^r = \sum_{j=1}^{N_r} \left[\theta_j^{sin.r} \sin\left(\frac{2\pi j}{P_r} t\right) + \theta_j^{cos.r} \cos\left(\frac{2\pi j}{P_r} t\right)\right] \qquad (6)$$

where $r \in \{m, y\}$ is the (monthly m or annual y) component indicator; P_r is the maximum cycle length for component r calculated as the number of days in a given month for $r = m$ and in a given year for $r = y$; $\theta_j^{sin.r}$ and $\theta_j^{cos.r}$ are the estimated parameters; and N_r is the maximum number of cycles used for component r.

Selecting relevant cyclical components is one of the problems of seasonality modelling in the parametric approach. On the one hand, their under-inclusion leads to cyclical fluctuations leaking into the cleaned series; on the other hand, too large a number can lead to overfitting. In order to address this problem, the model includes a large number of components ($N_m = 10$, $N_y = 20$), and to prevent overfitting at the estimation stage, Bayesian regression with automatic selection of hyperparameters is used (Tipping (2001)) that drops irrelevant components by maximising the marginal likelihood. Although marginal likelihood can be calculated directly for Bayesian regression, we follow the Mean Field (MF) approximation estimation procedure described in Khabibullin and Seleznev (2020), since later on, it will allow us to estimate models with non-linearities and non-deterministic trends with little modification.

Extensions of the basic procedure

In some situations, the basic model may not be flexible enough to adequately isolate the seasonal component (e.g. if 2020 is included in the sample) or exclude an insufficient number of patterns for subsequent analysis. In these cases, it can be easily modified by adding extra patterns. The model described in the previous subsection was chosen for a specific purpose of cleaning financial flows data of a specific set of patterns. However, the estimation methodology is flexible enough and can be easily applied without much modification to any model where the likelihood function can be expressed as a function of parameters and hidden variables.

Properties of the basic procedure

As we mentioned in the introduction, the choice of the seasonal smoothing procedure was dictated by a number of factors, including the quality of pattern removal discussed in the next section, as well as some features of the data and the need to address several technical difficulties. In this subsection, we will discuss how the proposed algorithms deal with these difficulties.

Despite the fact that many algorithms (for example, Facebook Prophet) can, in their original form, or with minor changes, cope with the problem of removing weekly, monthly and annual seasonality, as far as we know, their standard implementations do not admit missing data. Our basic procedure, which is essentially a Bayesian regression, can easily deal with missing data because it allows us to omit it from the sample. If the modifications contain non-deterministic components, the procedure is written in the form of a state-space model, where there are no observations in the specified time periods.

As mentioned above, the variational Bayesian inference using the MF approximation can be easily applied to a wide range of models. It also requires literally a few lines of additional code when making changes to the model[4], which allows easy modifications to the basic procedure. Changing other algorithms requires a lot of time and effort either to develop new estimation algorithms or to make edits to the program code of standard libraries, which is often quite laborious.

Finally, for four years of data (more than 1000 points), the running time[5] of the basic procedure does not exceed 2 minutes, and for extensions using nonlinear state space models does not exceed 10 minutes, which, if necessary, allows daily reevaluation of models for all industries[6] and testing new specifications in a reasonable time.

3. Quality criteria

Seasonal component extraction is an example of a problem where the correct answer is unknown, and where solution for this problem does not rely on training samples. For such problems, it is not possible to estimate the quality of their solution in a classical way by calculating a loss function on a test sample; therefore, we need a set of indirect quality criteria to assess whether the task is performed well. When removing seasonality from financial flows data, we rely on the quality of the forecast, on metrics that allow us to assess the presence/absence of overfitting and underfitting, as well as on a visual assessment.

Forecast quality. Although predictive properties do not always correlate well with the procedure for cleaning data of the seasonal component, a big difference in forecasts

[4] We use Tensorflow library (Abadi et al. (2016)).
[5] Hardware characteristics: Intel(R) Core(TM) i7-8750H CPU @ 2.20GHz 2.21GHz, RAM 16 GB.
[6] In practice, the coefficients of the seasonal component are reassessed no more than once a week.

between the seasonal adjustment model and its alternatives may be an evidence that the model is quite misspecified, and estimates are far from the desired result.

Lack of overfitting. Overfitting leads to situations when the model removes certain patterns that are not actually present in the data generating process in addition to the patterns that it should remove since the model has too much variation. To test this behaviour, we perform standard test-training diagnostics (Ng (2019)).

Lack of underfitting. Underfitting leads to the opposite situation. This happens when the model specification is not flexible enough or is overly regularised. To detect this, we build a series of additional regressions of model residuals on cyclic components that describe fluctuations of various frequencies.

Visual assessment. Human judgment about how well the model solves the problem is an important criterion for many tasks where direct quality measurement is not possible or standard metrics cannot fully capture all the nuances.[7] In this paper, we perform a visual comparison of the isolated seasonality and the original series, as well as an analysis of the cleaned series.

4. Sectoral financial flows data

The data represent firms' incoming[8] daily payments via the BRPS, aggregated by industry according to each firm's main code of activity in the Russian industry classification (OKVED2).[9] In addition to industry-specific payments, we have also isolated payments received by individuals and other payments, which are further indicated by codes 0 and 100, respectively. Time series are available from 1 January 2016 to the present[10] and have certain specific features that, as we have mentioned in the introduction, influenced the choice of the seasonal adjustment methodology. The main feature is that there are days when no payments were made. This is characteristic of holidays and weekends; however, due to inaccuracies in filling out payment orders, several payments still fell on weekends. These payments were excluded from our review to ensure the robustness of the data cleaning procedure.

5. Results

The results of the forecasts obtained using the basic procedure were compared with several widely used seasonal adjustment models: Facebook Prophet (Taylor and Letham (2017)), Seasonal and Trend decomposition using Loess (STL, Cleveland et al. (1990)), and Trigonometric Box-Cox transformation ARMA Trend Seasonality (TBATS, Livera et al. (2011)).[11] To the best of our knowledge, none of the software packages implementing these

[7] This practice is often used when building generative image (Arjovsky et al (2017)) or language models (Brown et al. (2020)).
[8] In this paper, we describe a methodology based on incoming payments. The same adjustment is applied to outgoing payments.
[9] Payments with the same taxpayer identification number (TIN) are excluded from the data.
[10] In this paper, the sample ends on 23 October 2020.
[11] Additionally, we estimated the Bundesbank seasonal adjustment model (Ollech (2018)), but its results were significantly worse than other alternatives, so we did not include them in Table 1.

methods can work with zero/missing data, and, therefore, we added logarithmic moving averages for 5 non-zero days. Table 1 shows the industry average RMSFE (relative to the linear trend model) for forecasting 2019 test sample at different horizons (the forecasting procedure is fully described in Appendix B). The SBL model works better, or is comparable to, other models,[12] which can serve as evidence that the seasonal component is extracted adequately. Table 1 also shows the RMSFE for linear regression (without any regularisation) using the same frequencies as in the SBL model (Regression) and using frequencies chosen on the basis of the economic considerations (Cut Regression).[13] Economically determined frequencies show slightly worse results than the basic procedure, while the regression results are about the same, which indirectly indicates that there is no overfitting.

The presence of overfitting is however indicated by the ratio of MSE in the test and training samples. Figure 2 shows that for many series this ratio is significantly greater than 1. This could be due to either overfitting or to the difference between the test and training samples. To understand the cause, we additionally compared the in-sample error ratio for the same forecasting periods in the model that was estimated on the data from 2016 to 2019. Figure 3 demonstrates a slight difference between the ratios of detrended errors for two different training periods and allows us to conclude that the test and training samples differ, rather than that overfitting is present. We run additional SBL regressions on residual series using the same set of trigonometric functions as in the original SBL model. As a result, we failed to find any evidence of underfitting.

Figures 4–9 show the seasonal components of incoming financial flows data for six different industries that reflect the typical results of our basic procedure. Figures 4–6 illustrate series with dominating daily, intramonthly, and intra-annual patterns, respectively. In all three cases, the SBL model captures the necessary frequencies and is sufficient from the point of view of visual analysis and the formal tests described above.

Although for most series, the properties of the extracted seasonal component are similar to those of the series in Figures 4–6, we also noticed a number of cases in which the basic procedure demonstrates unsatisfactory behaviour for subsequent analysis and requires additional adjustment.[14] These situations are shown in Figures 7–9 and can be addressed by applying various modifications. The first one is related to the changed behaviour of the series and the resulting changes in seasonality. We did not study this series further due to the obvious structural shift in the behaviour of seasonality that does not match the definition of a periodically recurring event and requires additional research.

The second situation is related to time-varying events. This situation is demonstrated in Figure 8. One can identify peaks concentrated around the 25th day of each month that cannot be fully explained by the seasonal component. This is due to the fact that the day on which the peak occurs in different months is located in a slightly different phase of the cycle,[15] which leads to a situation, in which seasonal component is averaged between peak and

---

[12] This may be because additional completion does not have a good effect on these procedures, or not enough time is spent on selecting hyperparameters; however, additional methods for correcting data and the validation procedure gives us hope that this is not the case.

[13] Quarterly fluctuations, monthly fluctuations, and the daily component.

[14] Currently, these adjustments are not applied in published reports, but shortcomings are taken into account when conducting the analysis.

[15] Both because the peak falls on different days of the month and because months have a different number of days.

ordinary days. We compared the residuals in the months where the 25th and 26th day of the month are business days and found that in only one out of 29 cases positive errors occur on both days. In other cases the sign of error differs (25^{th} day of the month – positive, 26^{th} – negative) with approximately the same in modulus mean error values: 0.68 and -0.7. Even though such patterns are not strictly periodic, their exclusion may be required for subsequent analysis. In each specific situation, the modification of the basic procedure should be considered and should depend on the purpose.[16] These modifications may include a set of dummy variables, a multinomial, or, in the case when the total payment is split over several days, the Dirichlet distribution. For the series in Figure 8, we add a set of dummy variables that are equal to one for the 25th day or the next business day, if the 25th falls on a weekend. Figure 10 shows that the seasonality removed in this way successfully copes with this type of pattern.

The third situation occurs when the amplitude of the series changes over time, as in Figure 9. To further remove such patterns, we include additional features in regression (2) in the form of the product of periodic features and time.[17] The results after adding new features are shown in Figure 11.

As described in Section 4, we only use data from 2016 to 2019. This is due to the presence of an abnormal period in the data after the introduction of restrictions caused by the spread of the coronavirus infection. To account for this behaviour and not limit ourselves to the pre-coronavirus period, we can modify the basic procedure by adding a flexible specification, such as a local linear trend model or a stochastic trend with stochastic volatility in trend innovations. To illustrate the possibility of using a model with strong nonlinearity, we demonstrate the second option. Figure 12 shows data that include an additional period up to 23 October 2020, as well as the series cleaned of seasonality using the basic procedure and the extended procedure with a flexible trend. We can see that the former is not able to take the drastic decline in incoming flows adequately into account and identifies false peaks in April. This is because to compensate for the low values of 2020, the model slightly overstates the values for previous years. In a model with a flexible trend, this does not happen because the April decline corresponds with the trend, which leads to superior visual results.

It is also worth noting that adding stochastic trends increases the running time of the algorithm about 5 times (from 2 to 10 minutes). However, a number of preliminary calculations for the local linear trend model showed that, despite slightly better forecast results, this does not lead to significant changes in the seasonal component on the pre-2020 data.

[16] We leave creating an automatic procedure for adding such patterns for future research. Currently, this is done manually for individual series. We also noticed that regressions of 1 to 5 day moving moduli of residuals on a set of sines and cosines are useful in detecting the need for these patterns.

[17] Alternatively, various data transformations can be used. However, after trying the Box-Cox transformation at the preliminary stage, we found that using stochastic optimisation algorithms for such models often leads to visually inferior results than simply adding new properties.

6. Related work

A large number of studies are devoted to estimating the seasonal component, but many of them, such as X-12, X-13 TRAMO-SEATS,[18] and JDemetra+[19] cannot be applied for our purpose, since they are designed to clean up monthly and quarterly time series.

Studies that describe seasonal adjustment algorithms for daily data can be divided into two categories: models with time-varying and constant seasonality. Even though some papers with time-varying seasonality, like ours, use trigonometric functions for constructing models (e.g. Livera et al. (2011), Ollech (2018)), they do not fit our definition of seasonality, as well as other models of this class that are based on local regressions and low-frequency filters (e.g. Cleveland et al. (1990), Verbesselt et al. (2010), and Wen et al. (2020)) or state-space models (e.g. Koopman et al. (2009), Koopman and Ooms (2006)). The second group is based on regression-type form of constant seasonal components. It includes the Prophet (Taylor and Letham (2017)), STR[20] (Dokumentov and Hyndman (2015)) algorithms and Campbell et al. (2005). These papers are the closest to our work and serve us as a starting point for constructing a procedure for seasonal adjustment of financial flows data. However, none of the implementations known to us can take into account all the specifics of our data (including missing values and flexible trends). As part of procedures that include RegARIMA models (Ghysels et al. (2001), Ollech (2018)), seasonality can also be modelled using seasonal lag polynomials. However, this requires intensive calculations due to a large number of periodic patterns that need to be considered.

Our work is also related to the area of research that deals with determining relevant frequencies. The most popular choice of the specification is based on information criteria (e.g. BAYSEA (Akaike (1980), Akaike and Ishiguro (1983)), as well as Taylor and Letham (2017), Ollech (2018) procedures). An alternative approach is to use regularisation with Ridge or LASSO regressions as in Dokumentov and Hyndman (2015). However, in contrast to the SBL procedure used in this work, these approaches require the model to be re-estimated for each specification, which can be computationally difficult and is not suitable in cases where estimating the model itself takes a long time.

Our basic procedure uses a linear trend and a normally distributed irregular component, which is close to the ideas applied in the Prophet library (Taylor and Letham (2017)) that uses a piecewise linear specification. However, extensions of the basic procedure can easily include any deterministic or stochastic model for both the trend and the irregular component, which is close in spirit to state-space models (e.g. Koopman et al. (2009), Koopman and Ooms (2006)).

[18] https://www.census.gov/srd/www/x13as/

[19] https://ec.europa.eu/eurostat/cros/content/software-jdemetra_en

[20] STR can also be estimated with varying seasonality.

7. Conclusion

In this paper, we have described the methodology for seasonal adjustment of daily data used by the Bank of Russia for preliminary cleaning sectoral financial flows. The simple basic procedure described in Section 4 removes recurring events well enough and can be easily modified to add new patterns and to use more flexible models if necessary.

References

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D.G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y. and X. Zheng (2016). Tensorflow: a System for Largescale Machine Learning, OSDI, 16, 265–283.

Akaike, H. 1980. 'Seasonal adjustment by a Bayesian modeling.' *Journal of time series analysis*, 1(1), 1–13.

Akaike, H., & Ishiguro, M. 1983. 'Comparative study of the x-ll and baysea procedures of seasonal adjustment.' *Applied Time Series Analysis of Economic Data*, 17–50.

Arjovsky, M., Chintala, S., & Bottou, L. 2017. 'Wasserstein gan.' arXiv preprint arXiv:1701.07875.

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... & Agarwal, S. 2020. 'Language models are few-shot learners.' arXiv preprint arXiv:2005.14165.

Carvalho, Vasco M. and Garcia, Juan R. and Hansen, Stephen and Ortiz, Alvaro and Rodrigo, Tomasa and Rodriguez Mora, Jose V. and Ruiz, Pep, 'Tracking the COVID-19 Crisis with High-Resolution Transaction Data,' *NBER Conference 'Economic Fluctuations and Growth Summer 2020' paper*, July 2020, Available at: http://conference.nber.org/conf_papers/f143494.pdf

Campbell, S. D., & Diebold, F. X. 2005. 'Weather forecasting for weather derivatives.' *Journal of the American Statistical Association*, 100(469), 6–16.

Cleveland, R. B., Cleveland, W. S., McRae, J. E., & Terpenning, I. 1990. 'STL: A seasonal-trend decomposition.' *Journal of Official Statistics*, 6(1), 3–73.

Chetty, R., Friedman, J., Hendren, N., & Stepner, M. 2020. 'The economic impacts of COVID-19: Evidence from a new public database built from private sector data.' Opportunity Insights.

De Livera, A. M., Hyndman, R. J., & Snyder, R. D. 2011. 'Forecasting time series with complex seasonal patterns using exponential smoothing.' *Journal of the American Statistical Association*, 106(496), 1513–1527.

Dokumentov, A., & Hyndman, R. J. 2020. STR: 'A seasonal-trend decomposition procedure based on regression.' arXiv preprint arXiv:2009.05894.

Ghysels, E., Osborn, D. R., & Sargent, T. J. 2001. *The econometric analysis of seasonal time series*. Cambridge University Press.

Hueng, C. James, ed. 2020. *Alternative Economic Indicators*. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.

Khabibullin, R., & Seleznev, S. 2020 ‘Stochastic Gradient Variational Bayes and Normalizing Flows for Estimating Macroeconomic Models.’ *Bank of Russia Working Papers Series* (Paper: wps61)

Koopman, S. J., & Ooms, M. 2006. ‘Forecasting daily time series using periodic unobserved components time series models.’ *Computational Statistics & Data Analysis*, 51(2), 885–903.

Koopman, S. J., Ooms, M., & Hindrayanto, I. 2009. ‘Periodic unobserved cycles in seasonal time series with an application to US unemployment.’ *Oxford Bulletin of Economics and Statistics*, 71(5), 683–713.

Lewis, Daniel J. and Mertens, Karel and Stock, James H. and Trivedi, Mihir 2020. ‘Measuring Real Activity Using a Weekly Economic Index.’ *FRB of New York Staff Report No. 920*, Revised September 2020, Available at: https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr920.pdf

Ng, A. 2019. ‘Machine learning yearning: Technical strategy for ai engineers in the era of deep learning.’ Retrieved online at https://www.mlyearning.org

Ollech, D. 2018. ‘Seasonal adjustment of daily time series’. *Discussion Paper No. 41/2018*. Deutsche Bundesbank.

Taylor, S. J., & Letham, B. 2017. ‘Forecasting at scale.’ PeerJ Preprints 5: e3190v2

Tipping, M. E. 2001. ‘Sparse Bayesian learning and the relevance vector machine.’ *Journal of Machine Learning Research*, 1(Jun), 211–244.

Verbesselt, J., Hyndman, R., Newnham, G., & Culvenor, D. 2010. ‘Detecting trend and seasonal changes in satellite image time series.’ *Remote sensing of Environment*, 114(1), 106–115.

Wen, Q., Zhang, Z., Li, Y., & Sun, L. 2020, August. Fast RobustSTL: ‘Efficient and Robust Seasonal-Trend Decomposition for Time Series with Complex Patterns.’ In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (pp. 2203–2213).

Appendix A. Figures and tables

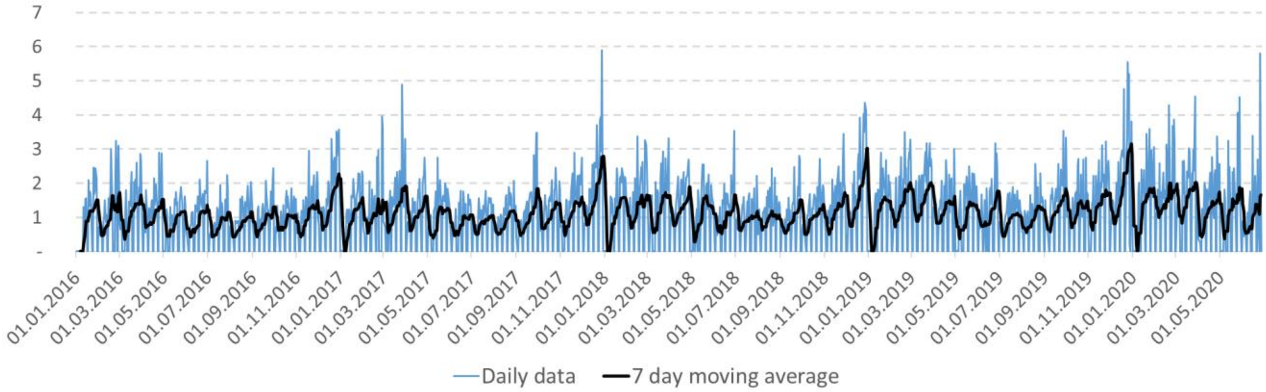

Figure 1. Production and distribution of electrical energy, gas, and water; incoming financial flows; normalised by the sample standard deviation.

	SBL	Regression	Cut Regression	Prophet no chngepoints	Prophet with changepoints	STL	TBATS
1 day	0,81	0,80	0,84	0,90	0,88	1,05	0,91
2 days	0,81	0,81	0,85	0,90	0,89	1,05	0,97
3 days	0,81	0,81	0,85	0,90	0,89	1,05	1,01
4 days	0,81	0,81	0,85	0,91	0,89	1,05	1,05
5 days	0,81	0,82	0,85	0,91	0,90	1,05	1,08
1 week	0,81	0,82	0,85	0,90	0,89	1,06	1,01
2 weeks	0,81	0,81	0,85	0,91	0,90	1,05	1,13
1 month	0,81	0,81	0,85	0,90	0,89	1,05	1,30
1 quarter	0,82	0,82	0,86	0,90	0,92	1,03	2,13

Table 1. Average relative RMSFE for 2019.

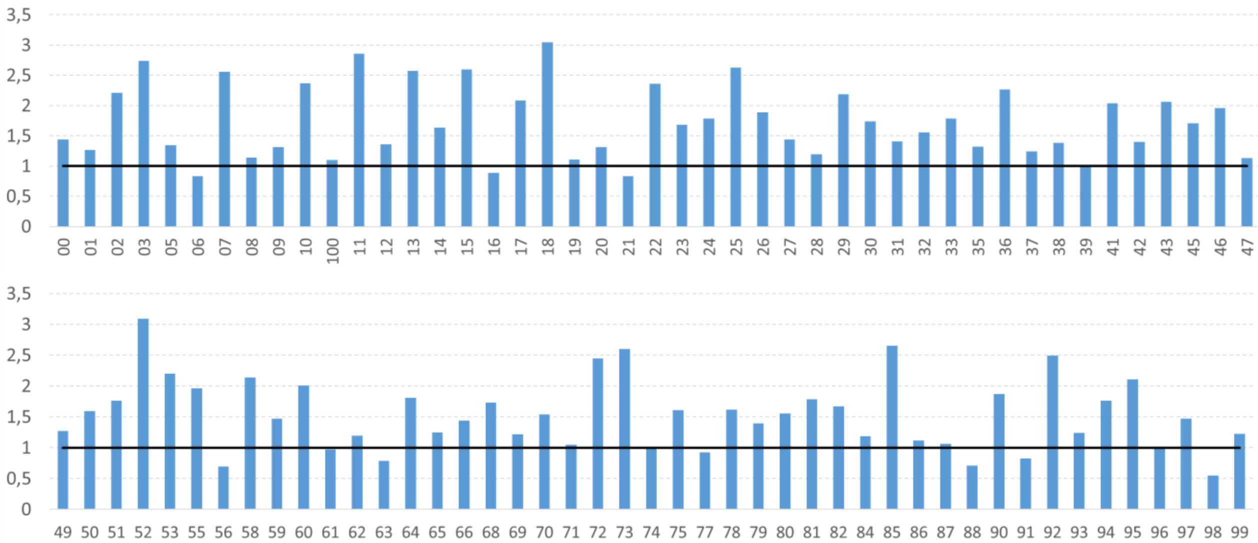

Figure 2. Ratio of MSE (1 day) for the test and training samples by sector; incoming flows.

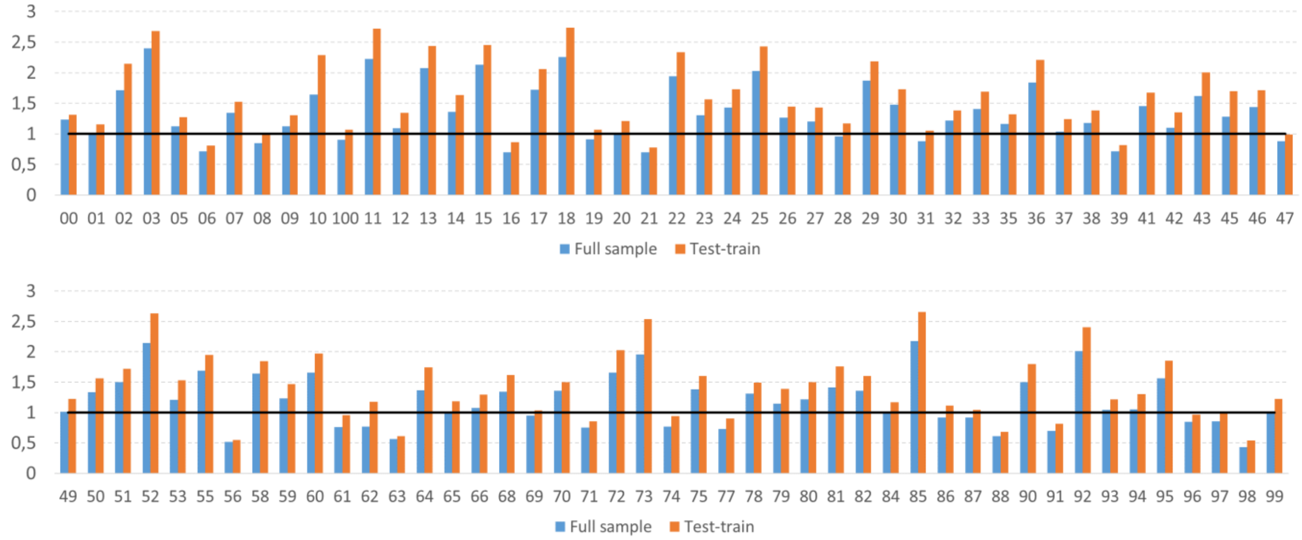

Figure 3. Ratio of mean squared detrended errors for two different training periods; incoming flows.

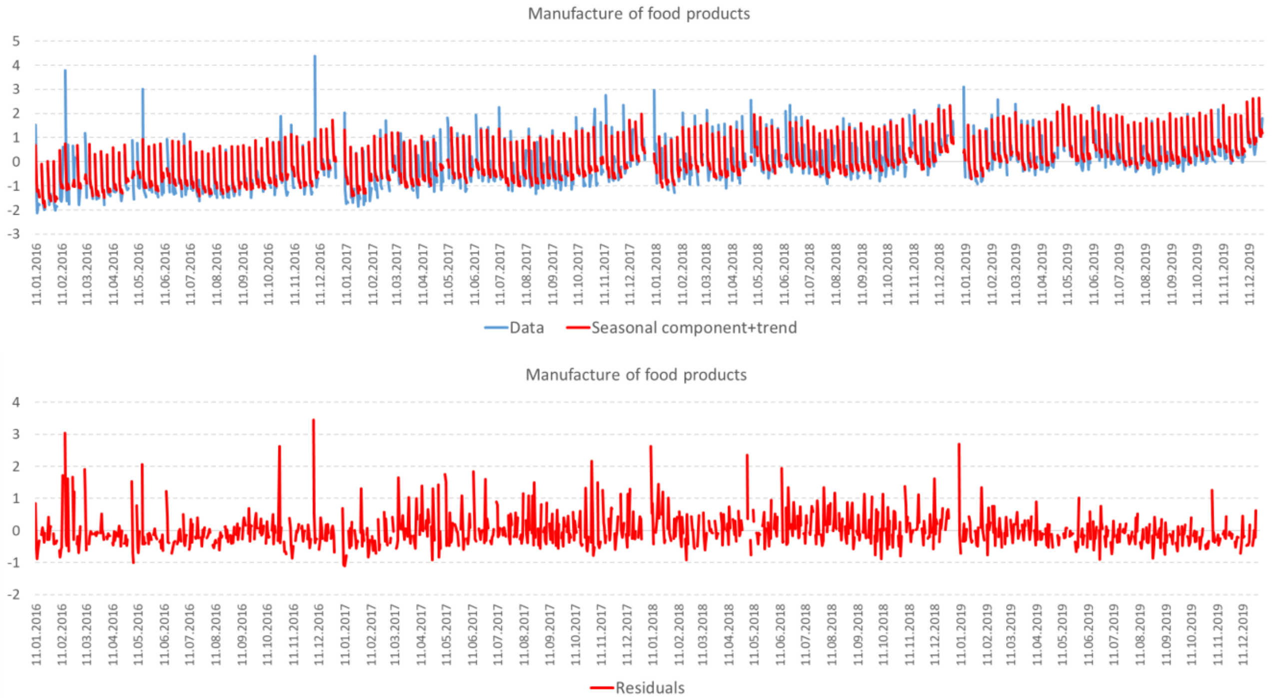

Figure 4. Seasonal component and residuals for manufacture of food products; normalised by the sample standard deviation.

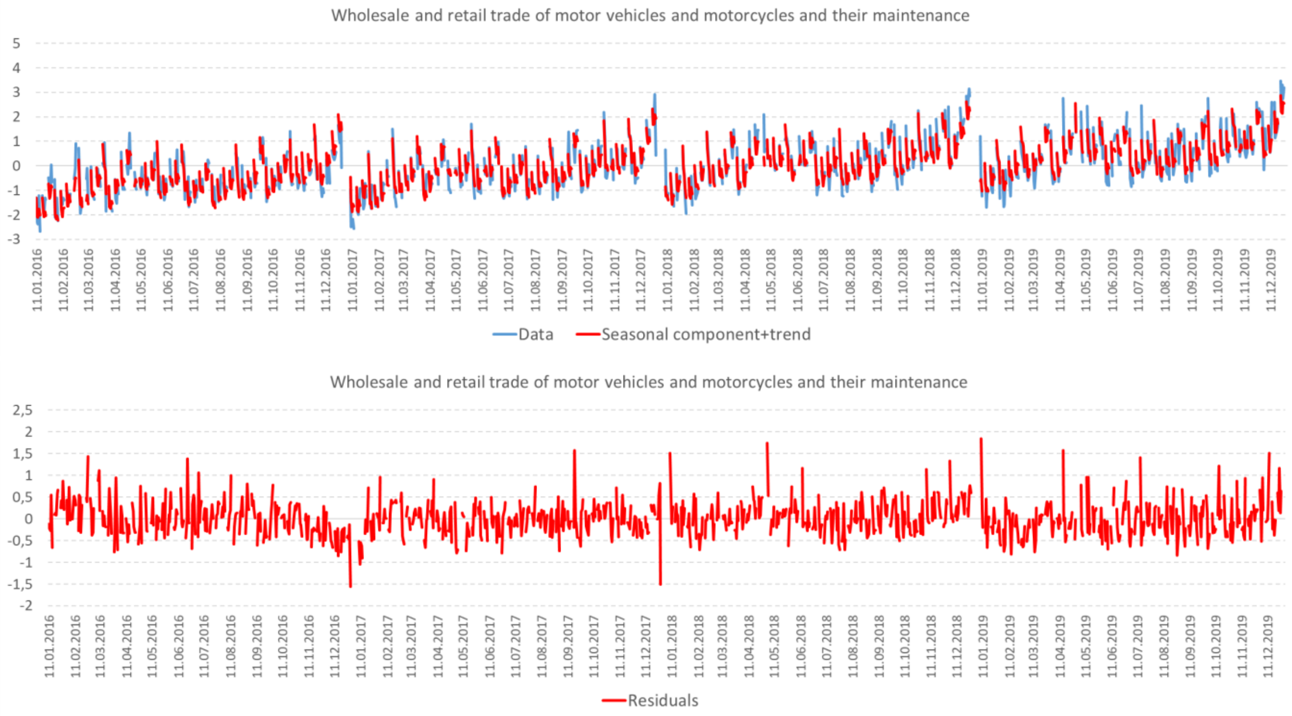

Figure 5. Seasonal component and residuals for wholesale and retail trade of motor vehicles and motorcycles and their maintenance; normalised by the sample standard deviation.

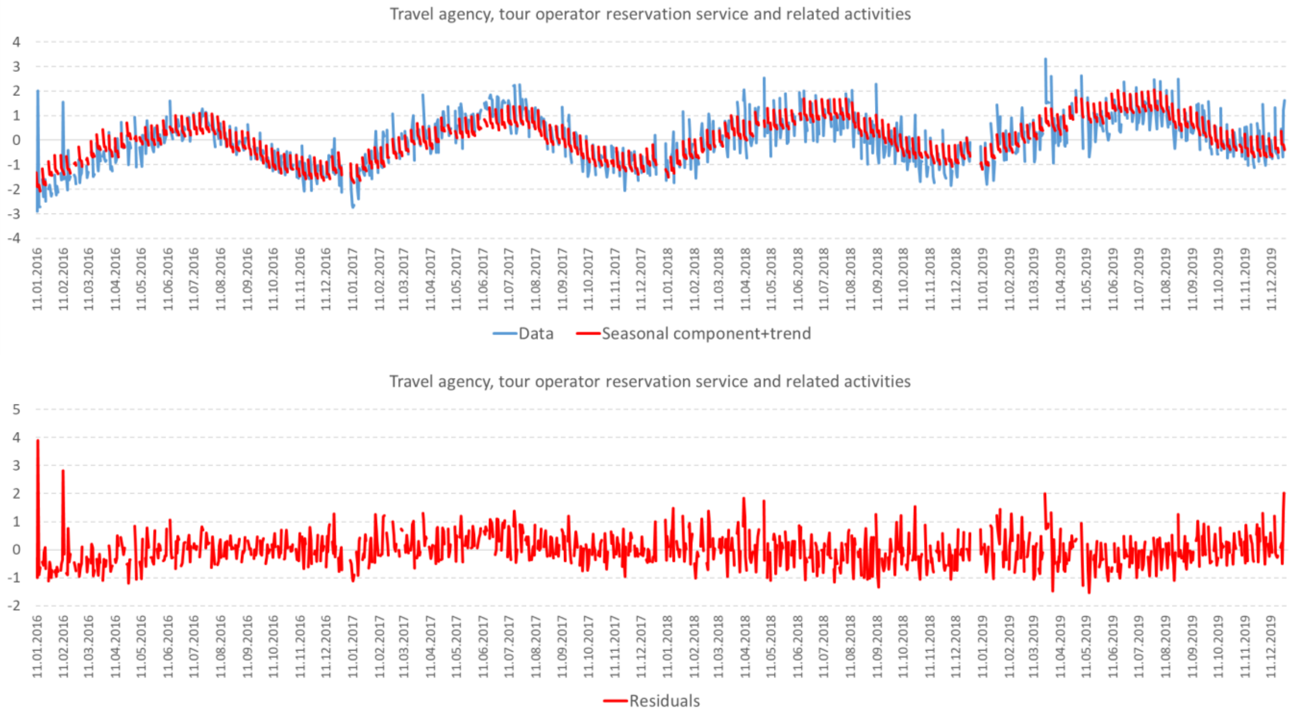

Figure 6. Seasonal component and residuals for travel agency, tour operator reservation service and related activities; normalised by the sample standard deviation.

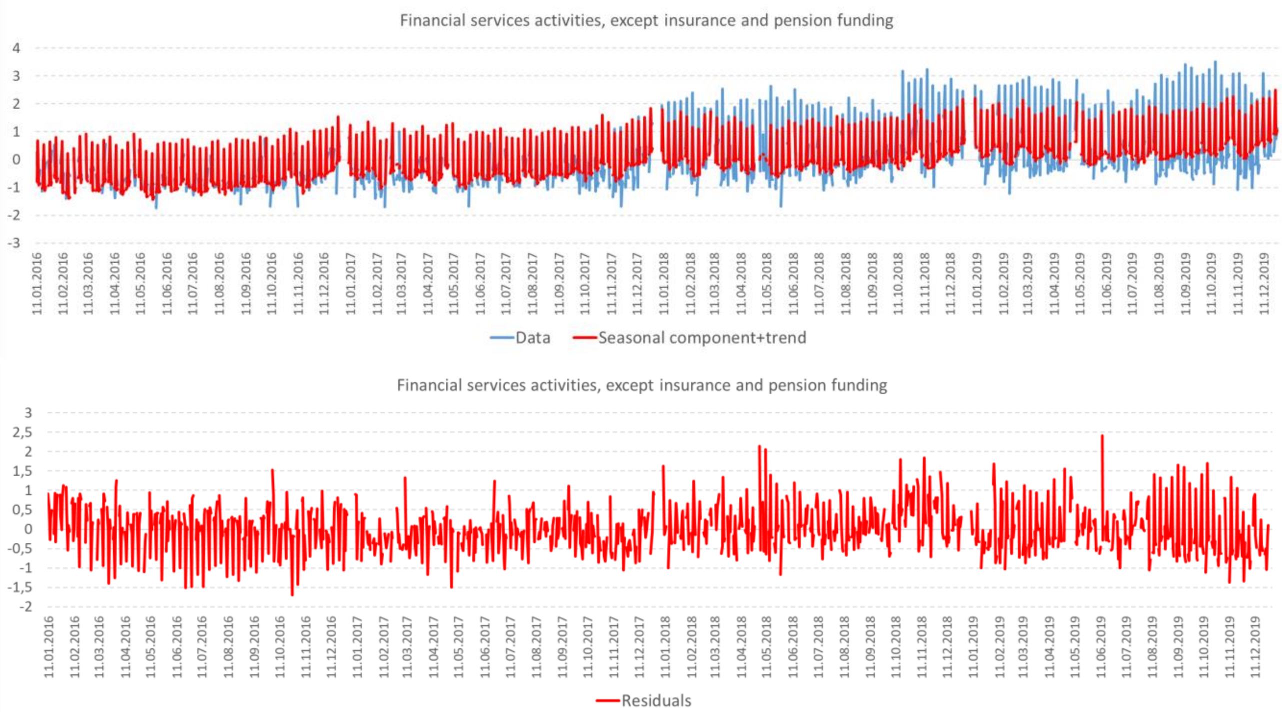

Figure 7. Seasonal component and residuals for financial services activities, except insurance and pension funding; normalised by the sample standard deviation.

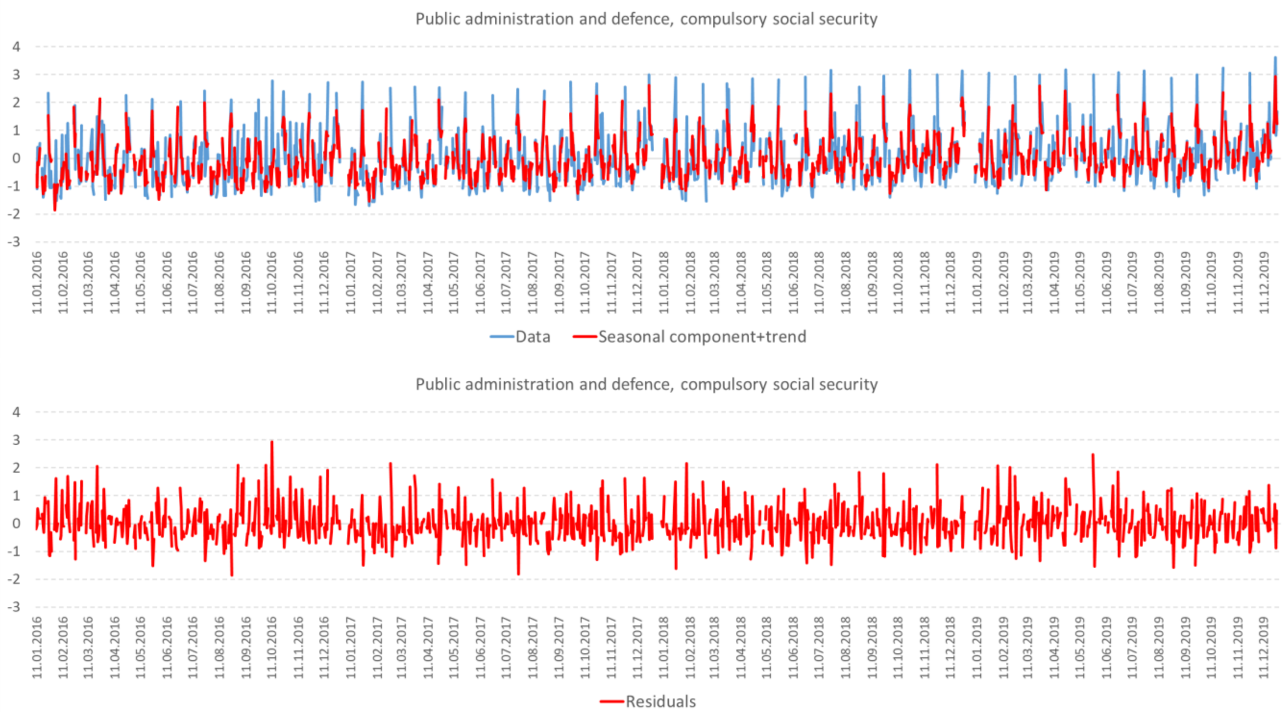

Figure 8. Seasonal component and residuals for public administration and defence, compulsory social security; normalised by the sample standard deviation.

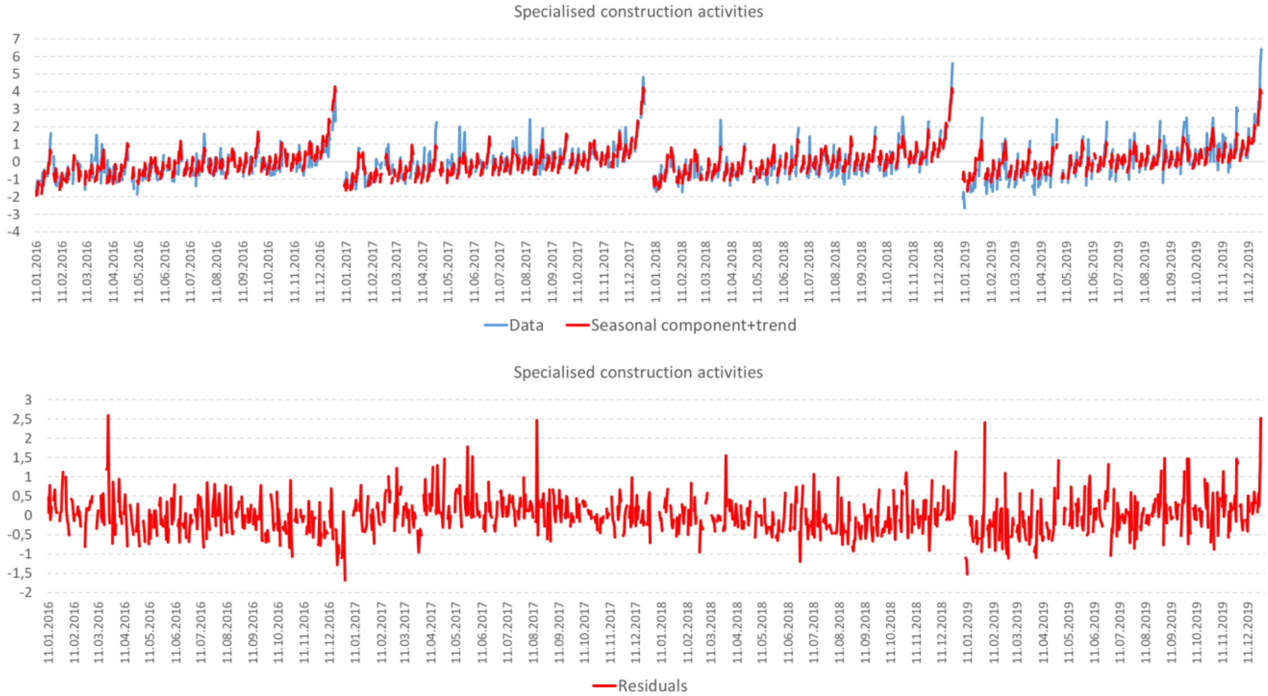

Figure 9. Seasonal component and residuals for specialised construction activities; normalised by the sample standard deviation.

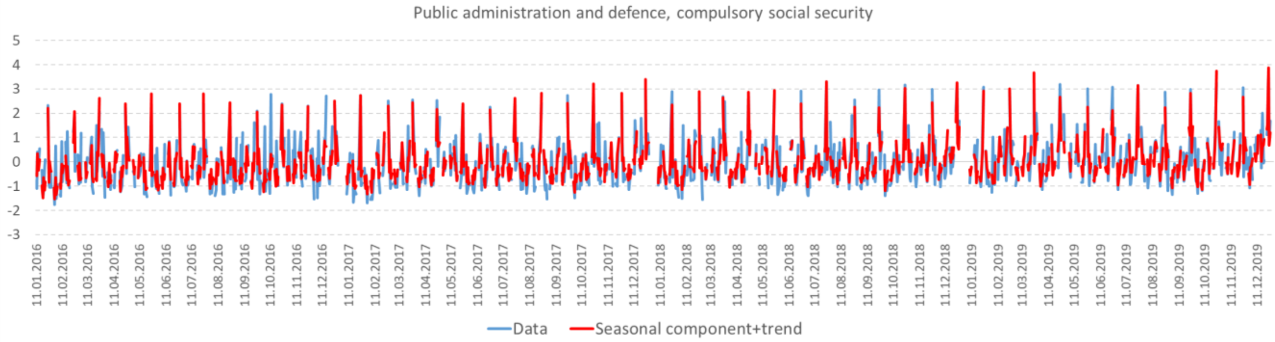

Figure 10. Adjusted seasonal component for public administration and defence, compulsory social security; normalised by the sample standard deviation.

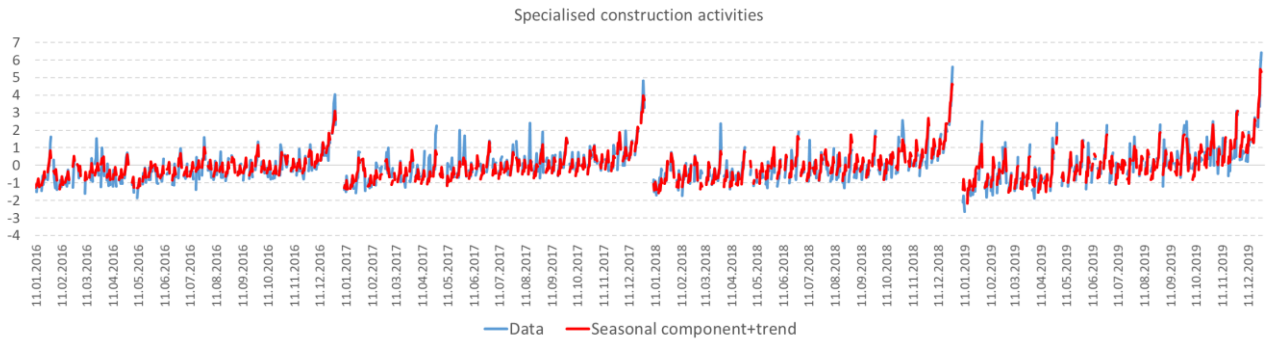

Figure 11. Adjusted seasonal component for specialised construction activities; normalised by the sample standard deviation.

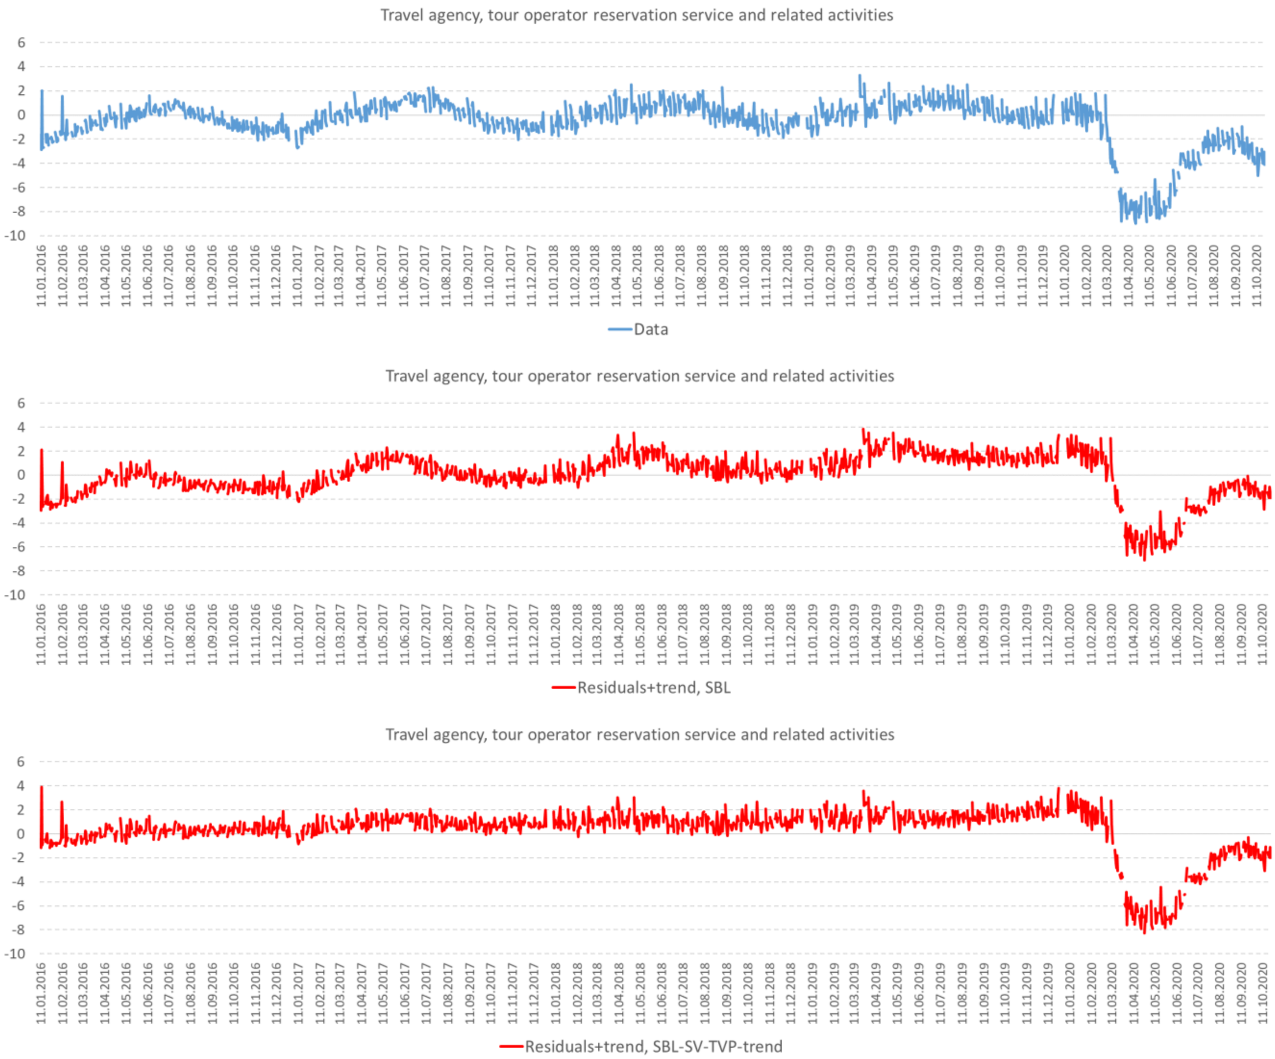

Figure 12. Data and residuals for travel agency, tour operator reservation service and related activities; normalised by the sample standard deviation for 2016–2019.

Appendix B. Forecasting procedure

In this paper, one of the criteria to determine the quality of seasonal component removal is the accuracy of its out-of-sample forecast calculated using RMSE for the logarithm of financial flows over several forecasting horizons. This includes a forecast for business days over the 1–5 days horizon, as well as forecasts for 7, 14, 28, and 84 calendar days.

To do this, each model was estimated on an expanding window using one-day increments on a sample consisting of all days of 2019, except for weekends and holidays. The first forecast was based on information available by the end of 2018 (28 December 2018), and the last forecast was based on information available by the penultimate point of 2019 (27 December 2019). For each forecast horizon $h \in \{1, \dots, 5, 7, 14, 28, 84\}$ days ahead a specific test sample was formed to estimate out-of-sample RMSE. For example, the test sample for $h = 1$ consisted of 246 points from 9 January to 30 December 2019, and for $h = 2$, from 10 January 2019 to 9 January 2020.

Holidays and weekends in the test sample were accounted differently for different forecasting horizons:

1. *Forecast for business days over the 1–5 days horizon.* Models were re-estimated on subsamples ending on business days. This was done recursively by sequentially adding one business day to the training sample. The forecasts were calculated for h business days ahead. For example, if $h = 2$ and the last day of the estimated sample was Friday, then the logarithms of the financial flows were forecasted for Tuesday of the next week. Holidays were excluded in the same way. If the sample ended right before the start of holidays, then for $h = 2$ second day after the holidays were forecasted.

2. *Forecasts for 7, 14, 28, and 84 calendar days ahead*, which approximately correspond to the forecast period of a week, two weeks, a month, and a quarter ahead. Models were also re-estimated recursively, adding one business day to the sample on which the model was estimated. However, the forecast was calculated for h calendar days (including holidays and weekends) ahead. At the same time, if the forecasted day turned out to be a weekend or holiday, this forecast was excluded from the test sample. Since the selected forecast horizons are multiples of 7, each time the same day of the week is forecasted, which means that no Saturday or Sunday can occur in the test sample. Thus, only forecasts for/on public holidays are excluded from the sample.

SBL regression is limited in that it was estimated only on a sample consisting of business days. However, the Prophet, STL, and TBATS methods were estimated on a sample consisting of calendar days, including weekends and holidays. In this case, all values for weekends and holidays were replaced with the average values of flows for the business week before the start of the corresponding weekend or holiday. However, for all forecast horizons, the test sample completely coincided with the SBL model. Moreover, each value of the test sample was forecasted from the same date as the SBL model.

Hyperparameters of models were estimated using the sample until the end of 2018 as if no information about the dynamics of variables in 2019 were available. However, hyperparameters estimation has its peculiarities for each model.

1. For SBL regression, prior variance hyperparameters were estimated on the sample from 1 January 2016 to 31 December 2018, excluding holidays and weekends.

2. For the Prophet models with and without change points, and for the STL model, a validation sample was constructed. Based on this sample an out-of-sample forecast was built similarly to the test sample, and RMSE were calculated. Model hyperparameters are selected based on the calculated RMSE. The validation sample was constructed in such a way as to select hyperparameters based on information up to the end of 2018 (28 December 2018), and at the same time so that the sample size was exactly one year. For example, if $h = 7$, the first forecast is based on the model estimate before 22 December 2017 (penultimate Friday of 2017), and the last forecast is based on the model using data prior to 21 December 2018 (penultimate Friday of 2018).

3. For the TBATS model, hyperparameters were selected based on the training sample, since the calculation using the validation sample is computationally complex. Forecasting using the test sample was implemented similarly to the Prophet and STL models.

For the LLT model (used in preliminary experiments), the forecasting procedure is different. All model parameters were estimated using data until the end of 2018. Then these parameters were used to calculate forecasts on the test sample. In order to predict the trend component using the expanding window, the forecast was estimated for each new sub-sample using the Kalman filter with fixed parameters. Regression coefficients for estimating the seasonal component were fixed in the same manner.