Seasonal Adjustment of Weekly Economic Series

Tim Ginker

Abstract

In recent decades, there has been a significant increase in the number of economic indicators available at high frequency. These indicators may reduce uncertainty and thereby assist policymakers in the decision-making process, especially during periods of rapid change that are typified by increased uncertainty. In order to understand the movements in these data, seasonal adjustment must first be carried out. However, seasonal adjustment of high-frequency data involves many challenges. This is due to the data's unique characteristics, which do not allow the application of standard methods commonly used by statistical bureaus worldwide, such as the X-13 ARIMA methodology. This paper presents the Bank of Israel's ongoing work in developing methodologies and statistical tools for seasonal adjustment of weekly frequency data. The use of open-source tools developed at the Bank is demonstrated on the basis of indicators of real economic activity and of the labor market in Israel.

1. Introduction

Policymakers' ability to track trends and identify turning points in economic activity is of paramount importance, especially during periods of rapid change that are typified by increased uncertainty. In recent decades, thanks to developments in information technology, there has been a significant increase in the number of high-frequency (i.e. weekly or daily) economic indicators that can reduce uncertainty and thereby assist in the decision-making process. However, the use of these indicators becomes more complex when they include not only a trend and irregular components but also seasonal effects. These effects do not represent the fundamental nature of the observed series, but rather technical factors such as the calendar (holidays, workdays, etc.) and other effects, the timing, direction, and magnitude of which remain relatively constant over time (such as weather conditions, students looking for jobs during semester breaks, etc.). As demonstrated below, it is not possible to infer the nature of changes in economic indicators without filtering out the seasonal component, which can even dominate the behavior of a series.

This working paper presents the Bank of Israel's ongoing activity in developing methodology and statistical tools for seasonal adjustment in weekly frequency data series. The use of open-source tools developed at the Bank will be demonstrated on the basis of indicators of real economic activity and of the labor market in Israel.

The first example will illustrate the seasonal adjustment of the weekly number of new registrations at the Israel Employment Service. This is administrative data available with less than a week's delay, which can provide an early indication of the unemployment situation in the economy. Figure 1 below shows the number of registrants at the Israel Employment Service, and clearly illustrates the difficulty in analyzing developments in the series without adjusting for seasonal effects. For example, two very sharp declines in the number of registrants are observed to recur every year. These declines do not reflect a genuine change in the unemployment situation in the economy but only the impact of moving holidays. There are also three sharp increases in the number of new registrants every year which are related to seasonal workers. The series is also characterized by strong intramonthly cyclicality, which further obscures the nature of the variation in the data.

Figure 1 – Number of registrants at the Israeli Employment Service

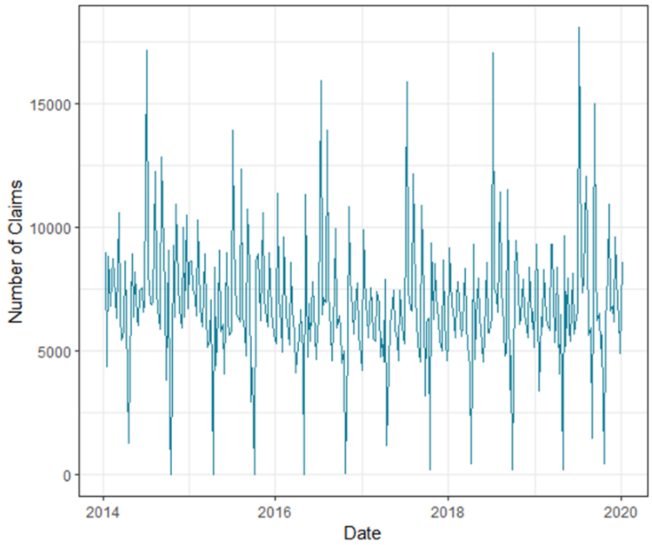

The second example relates to credit card purchases, which may provide an early indication of changes in private consumption. Unlike National Accounts data, which are published with a delay of about six weeks, daily and weekly data on credit card purchases are available in real time. Figure 2 shows weekly credit card purchases. This series exhibits seasonal cycles and effects arising from workdays and holidays, which need to be filtered out before analyzing the data.

Figure 2 – Credit card purchases

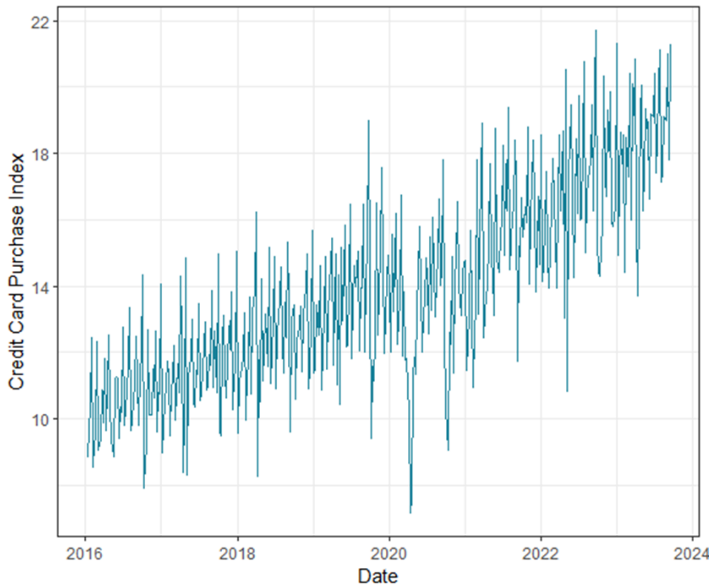

1.1 Challenges in adjusting high-frequency data

Seasonal adjustment of high-frequency data involves numerous challenges. This is due to the data's unique characteristics, which do not allow for the general application of standard methods commonly used by statistical bureaus worldwide, such as X-13 ARIMA. In addition to preadjustments[1] and handling of outliers, most seasonal adjustment methods include two main steps. In the first, the trend component is estimated and removed, and in the second, the seasonal factors are estimated by smoothing the subseries within the seasonal cycle[2]. For example, after removing the trend in monthly data, the seasonal factors are estimated by calculating the (weighted) average of each month separately.

This method of seasonal adjustment is generally not applicable in the case of weekly series because the length of the cycle may vary from period to period, there may be multiple seasonal cycles, and the position of the week within the month may shift. For instance, the number of weeks in a year can be 52 or 53. Additionally, weekly data can include several seasonal cycles, such as both the intramonthly and intrayearly cycles. This can be illustrated using the employment data presented in Figure 1. Suppose that the number of weeks is the same every year, as well as the number of weeks in a month, making it possible to apply the smoothing to subseries, that is, to calculate the average number of registrants by the position of the week in the month (for example, the average of the first to the fourth week of each month) and also the average number of registrants according to the number of weeks in the year. This calculation will yield a biased result because of the location of the weeks within the month in a different part of the annual cycle. Furthermore, some weeks might even fall during the sharp increases that belong to the annual cycle that we saw earlier, illustrating the potential magnitude of the bias.

Another challenge is that the position of the weeks within the monthly cycle is not fixed. To illustrate, consider the example in Figure 3, which refers to the last three months of 2023. We created a deterministic daily series (black line) with a constant level and an intramonthly seasonal cycle, so that the series has high values at the beginning of the month, which gradually decrease linearly. Before we proceed with the example, it is important to note that because the increases and decreases are purely technical, we would expect a good seasonal adjustment method to eliminate the intramonthly cyclicality and to generate a horizontal linear line.

We now examine the characteristics of the series when observed at weekly frequency (as in the Employment Service data mentioned earlier), where the aggregation rule is such that each observation represents the average daily value over a week. Without loss of generality, we will focus on the first week of each month. The rectangle on the left side of the Figure shows that the first week of October represents the average of the beginning of the intramonthly cycle, which is marked by a green triangle. In contrast, the first week of November, which falls between October

[1] The preadjustment factors include the impact of Jewish holidays, the timing of which is not fixed in the Gregorian calendar, as well as the impact of The preadjustment factors include the impact of Jewish holidays, the timing of which is not fixed in the Gregorian calendar, as well as the impact of business days and workdays..

[2] We assume that the seasonal cycle includes *n* observations. For example, in monthly data, the annual seasonal cycle is 12. We define the series of observations located at a specific position in the seasonal cycle as a "cycle subseries." For example, in monthly data, the first cycle subseries of the seasonal cycle will include all Januarys, the second will include all Februarys, and so on.

29 and November 4, represents the aggregation of a different part of the intramonthly cycle, which is also true for the first week of December. As a result, due to the different positions within the monthly cycle, it is meaningless to use methods based on the smoothing of subseries within the cycle, which essentially assume the same seasonal effect in the first week of each month. It is worth noting that the same shift occurs in the intrayearly cycle (if it exists in the series).

Figure 3 – An example of the shift of intramonthly cycles in the weekly window

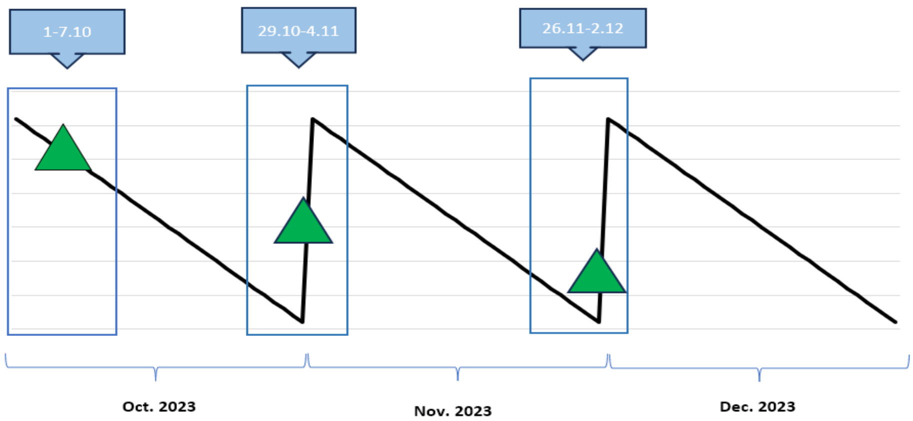

1.2 A review of methods and the literature

Although the first attempts to adjust for seasonality in weekly series began in the early 20th century (Crum, 1927), the literature in this area remains relatively limited. Similarly, statistical software for weekly seasonal adjustment is less developed and less accessible than tools designed for monthly or quarterly series.

Currently, there are several statistical tools for weekly seasonal adjustment in various stages of development. One example is the MoveReg software developed by the US Bureau of Labor Statistics, which can be operated using statistical tools like SAS and Eviews, and its methodology is well-documented by Cleveland et al. (2014). However, there are hardly any open-source tools, and even the most promising ones, such as the Ecce Signum R package (McElroy and Livsey, 2022), are still in development.

In addition to the methodologies described above, there are several other methods occasionally used for weekly seasonal adjustment. These methods were first proposed by Cleveland et al. (1990), and are based on seasonal-trend decomposition using LOESS (STL; where LOESS stands for locally estimated scatterplot smoothing). Recently, Bandara et al. (2021) extended the STL method to a multiple seasonal-trend decomposition using LOESS (MSTL), allowing for the decomposition of series with multiple nested seasonal cycles. It is important to note that both methods are based on the smoothing of cycle subseries, which is not suitable for weekly series due to the limitations described above. Additionally, the inability of both methods to incorporate preadjustment factors (such as holidays and workdays) likely explains why these methods have not been adopted by statistical bureaus worldwide.

There are several studies that suggest the use of forecasting models for seasonal adjustment, such as Prophet developed by Taylor and Letham (2018) or TBATS, which was proposed by De Livera et al. (2011). However, the findings in the literature indicate that the accuracy of the decomposition performed by forecasting methods is insufficient, likely due to a different objective function and reliance solely on past observations for calculating trend and seasonal factors (Bandara et al., 2021).

The rest of this working paper is structured as follows: Section 2 describes a methodology for seasonal adjustment of weekly data. Section 3 provides several examples of using the boiwsa[3] package developed at the Bank of Israel for adjusting series like the weekly number of registrants at the Employment Service and a weekly credit card purchase index. Section 4 summarizes the results of the working paper, offers suggestions for improvement, and outlines directions for future research.

[3] The package is available from CRAN. See https://cran.r-project.org/package=boiwsa for further details.

2. Methodology

This section presents our methodology for the seasonal adjustment of weekly data, which has been implemented using the boiwsa package. The method is based on the model developed by Cleveland et al. (2014), with a few differences designed to introduce greater flexibility and to facilitate the use of the tool for automatic adjustment.

We assume that the observed series y_t can be decomposed as follows:

$$y_t = T_t + S_t + H_t + O_t + I_t, \qquad (1)$$

where Tt represents the trend component, St the seasonal component, Ht the holiday and trading-day effects, Ot and It the outlier and irregular components respectively, and t denotes the date of the last day within a given weeknot.

The seasonal component is modeled using trigonometric variables as follows:

$$S_t = \sum_{k=1}^{K}\left(\alpha_k^y \sin\left(\frac{2\pi k D_t^y}{n_t^y}\right) + \beta_k^y \cos\left(\frac{2\pi k D_t^y}{n_t^y}\right)\right) + \sum_{l=1}^{L}\left(\alpha_l^m \sin\left(\frac{2\pi l D_t^m}{n_t^m}\right) + \beta_l^m \cos\left(\frac{2\pi l D_t^m}{n_t^m}\right)\right), \qquad (2)$$

where D_t^y and D_t^m represent the day in the year and the day in the month of observation t, and n_t^m and n_t^y represent the number of days in the corresponding month and year. Therefore, the model allows for the estimation of two seasonal cycles: the intrayearly and the intramonthly. To allow the seasonal component to vary over time, the coefficients in Equation (2) are estimated using a variant of Discount Weighted Regression (Harrison and Johnston, 1984). For more details on the implementation of this method in the boiwsa package, see Ginker (2023).

As in the case of the X-11 method (Ladiray and Quenneville, 2001), our procedure is based a recursive process to generate its various components as illustrated in the following chart:

Chart 1 – The stages in the seasonal adjustment process

Estimation of the trend component - T_t

↓

Calculation of the detrended series $y_t - T_t$

↓

Optional stage for the automatic calculation of the optimal number of variables in Equation (2) and a search for outliers

↓

Calculation of the seasonal component (including the effect of holidays and workdays, if needed)(

↓

Calculation of the seasonally adjusted series $y_t^{SA} = y_t - S_t - H_t$.

↓

Calculation of the trend within the adjusted series and after the neutralization of outliers

The process described above is repeated twice, where in the second iteration we use the estimated trend generated by the first iteration instead of reestimating the trend.

3. Empirical Application

In this chapter, we present two examples of seasonal adjustment of data available at weekly frequency. We also discuss several practical aspects related to the unique characteristics of the series and ways to handle them. The first example will demonstrate seasonal adjustment of the weekly number of new registrants at the Israel Employment Service, while the second will address weekly credit card purchase data.

3.1 Seasonal adjustment of the number of new registrants at the Israel Employment Service

The first example is based on the number of new registrants at the Israel Employment Service, as shown in Figure 1. The data cover the period from 2014 to 2020. The series is characterized by two nested seasonal cycles, as well as the nonstandard impact of moving holidays and the effect of workdays.

Every individual wishing to claim unemployment benefits must register with the Employment Service. For this reason, registration is expected to occur shortly after job termination. Given that most employment contracts end at the end of a calendar month, we expect to see more registrations at the beginning of the month, which will be reflected in an intramonthly seasonal cycle. Additionally, Figure 1 shows several recurring peaks each year, with the last one occurring toward the end of August. These peaks are related to seasonal workers, and form the basis of the intrayearly cycle.

Additionally, every year there are weeks when the number of registrants drops to almost zero. These declines are related to the impact of two moving holidays (Rosh Hashanah and Passover). In this case, preadjustment factors for moving holidays are not produced using standard methods (see the Easter formula in Table 2 of Findley et al. (1998)), because there is no distribution of the impact between weeks, but rather a decrease and an increase in the adjacent weeks. Therefore, these effects were adjusted using globally centered dummy variables (which can be reproduced using the function boiwsa::my_rosh). In addition to moving holidays, we also expect the number of workdays in a week to affect the number of registrants, with fewer full workdays resulting in fewer registrations. We adjust for these effects using unique preadjustment variables calculated by the function boiwsa::simple_td. For detail s on the implementation of the process in the boiwsa package, see Ginker (2023).

The analysis in Figure 4 shows that the procedure successfully eliminated the effects of two seasonal cycles combined with the impact of moving holidays. However, there is a noticeable sharp decline in the adjusted series starting from 2016. This decline recurs in subsequent years with decreasing intensity.

Figure 4 - The number of registrants at the Israel Employment Service in contrast to the adjusted series (in black)

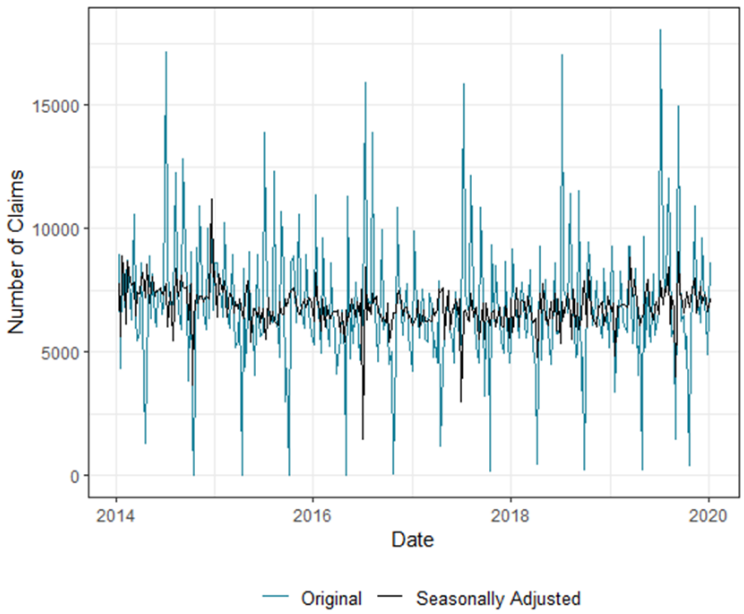

As mentioned above, the higher the frequency at which data is observed, the greater its volatility typically is. This characteristic makes it difficult to identify outliers and can distort the estimated seasonal factors, both in cases of incorrect identification of an observation as an outlier and when an outlier is not properly addressed. The impact of this distortion will weaken due to the local structure of the model, as shown in Figure 4. In this application, the adjustment relied on the package's automatic process to select the appropriate model and identify outliers (implemented using the function boiwsa::find_outliers), with identification based on the threshold values proposed by Findley et al. (1998) for monthly data, which might not be conservative enough for higher-frequency data. Figure 5 shows that the problem is resolved by raising the t-statistic threshold for the outlier variables. (For more details on setting the threshold values, see the package documentation.) It is worth noting that raising the threshold might come at the cost of failing to identify some outliers. This example underscores the importance of experts examining the results of the adjustment as well as judicious intervention when necessary.

Figure 5 – Number of registrants at the Israel Employment Service compared to the adjusted series (in black)

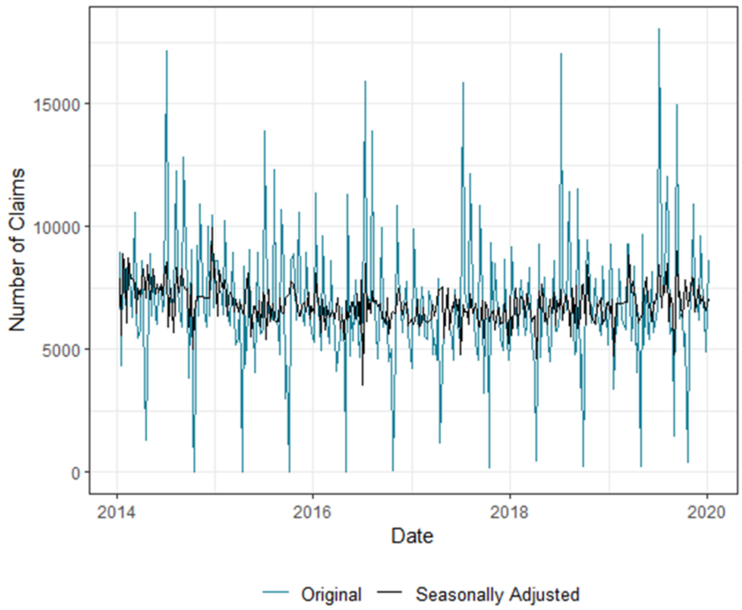

After the removal of the trend component[4], the quality of the adjustment can be assessed by comparing the spectrum[5] of the original series to that of the adjusted series. This procedure is performed using the function boiwsa::plot spec. To facilitate the analysis, we marked the first peaks belonging to the intramonthly and intrayearly cycles in the chart produced by the package. As shown in Figure 6, and in line with our initial hypothesis, the series is characterized by two seasonal cycles, which were successfully removed by the procedure.

[4] In spectral analysis, it is customary to remove the trend component before the analysis in order to prevent it from concealing the peaks related to the seasonal component. If this is not done, then in a series with a strong trend component, we will mainly see peaks at the start of the axes, which will make it difficult to identify the other peaks.

[5] The spectrum of a series is used to describe its variation as a function of cycles at different frequencies. Thus, significantly high values at the frequencies marked in the chart may indicate the presence of a seasonal cycle, and their absence in the adjusted series indicates the success of the adjustment procedure.

Figure 6 – The spectrum of the number of registrants at the Israel Employment Service

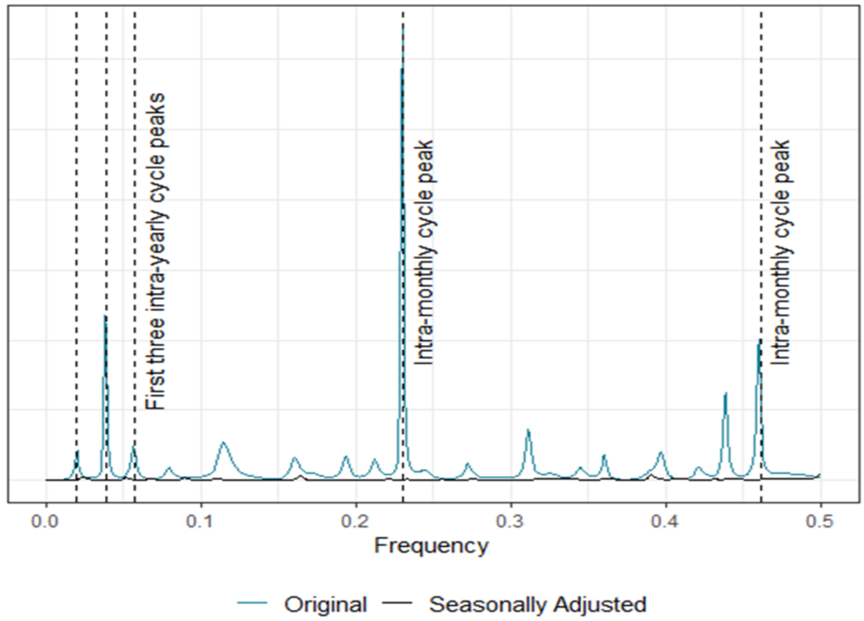

3.2 Adjusting the data for credit card purchases

The second example we present is based on the weekly Credit Card Purchase Index. These data are available daily and reported to the public by the Bank of Israel.[6] For this application, we converted the unadjusted series to a weekly frequency by summation. The data cover the period from January 9, 2016 to September 16, 2023, when this working paper was written. As with the data in the previous example, we anticipate a combination of several seasonal cycles as well as the influence of holidays and workdays. Factors such as the timing of salary transfers to bank accounts and purchasing habits may create an intramonthly seasonal cycle. There may also be effects due to workdays and moving holidays, and an intrayearly cyclicity arising from seasonal differences in purchasing habits over the course of the year.

Figure 7 shows the original series (in blue), with the part affected by moving holidays highlighted in grey, alongside the adjusted series (in black). In this application, preadjustment factors for moving holidays were produced using the Easter formula (see Table 2 of Findley et al. (1998)) using the function boiwsa::genhol, which is suited to weekly data. The package also includes a file with holiday dates, and another containing the classification of each day as a full workday, a partial workday, or holiday (see boiwsa::dates_il). As the graph shows, the procedure successfully adjusted for the effect of moving holidays. The variable for workdays was also significant, and its effect was removed from the series.

[6] See https://edge.boi.gov.il/#?locale=he.

Figure 7 – The Credit Card Purchase Index compared to the adjusted series (in black)

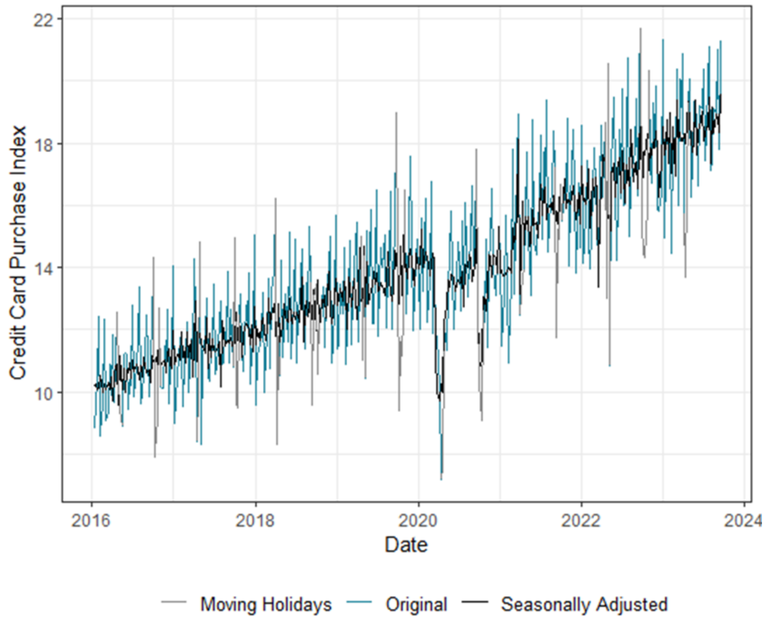

In addition, the spectrum analysis in Figure 8 shows that all seasonal peaks have been removed, with the most significant belonging to the intramonthly cyclicity.

Figure 8 – Spectrum of credit card purchase data

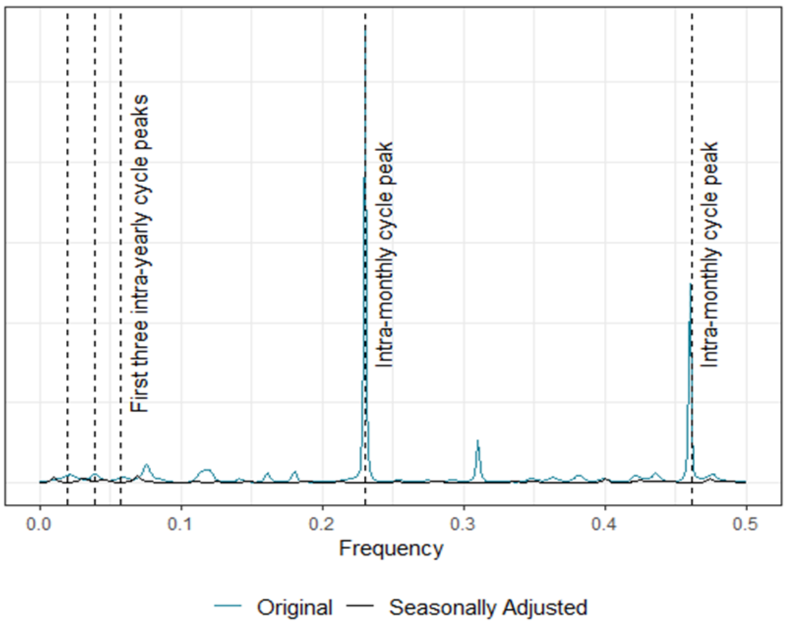

4.Conclusion

In this working paper, we have presented the ongoing work at the Bank of Israel to develop methodology and statistical tools for seasonal adjustment of weekly data. One of the directions for the future is to examine additional methods for trend extraction and the local structure of regression. For example, in the current version, the update rate of the intramonthly and intrayearly seasonal factors is identical, which may be limiting in certain applications and when extending the use of the method to higher-frequency data. The method can be made more flexible by setting different update rates. We also intend to expand the functionality of the boiwsa package, including the ability to perform adjustments according to additional methods.

5.Bibliography

Bandara, K., Hyndman, R. J. and C. Bergmeir (2021). MSTL: A seasonal-trend decomposition algorithm for time series with multiple seasonal patterns. arXiv preprint arXiv:2107.13462 .

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

Cleveland, W. P., Evans, T. D. and S. Scott (2014). Weekly Seasonal Adjustment - A Locally weighted regression approach. Economic working papers, Bureau of Labor Statistics.

Crum, W. L. (1927). Weekly fluctuations in outside bank debits. *Review of Economics and Statistics 9*(1), 30–36.

De Livera, A. M., Hyndman, R. J. 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.

Findley, D. F., Monsell, B. C., Bell, W. R., Otto, M. C. and B. C. Chen (1998). New capabilities and methods of the X-12-arima seasonal-adjustment program. *Journal of Business & Economic Statistics 16*(2), 127–152.

Ginker, T. (2023). Seasonal Adjustment of Weekly Data by Discounted Least Squares in R. *SSRN Electronic Journal*.

Harrison, P. J. and F. R. Johnston (1984). Discount weighted regression. *Journal of the Operational Research Society 35*(10), 923–932.

Ladiray, D. and B. Quenneville (2001). Seasonal adjustment with the X-11 method. Springer New York, NY.

McElroy, T. S. and J.A. Livsey (2022). Ecce Signum: An R Package for multivariate signal extraction and time series analysis. arXiv preprint arXiv:2201.02148 .

Taylor, S. J. and B. Letham (2018). Forecasting at scale. *American Statistician* *72*(1), 37–45.