Financial Risk&Regulation

Treatment of financial time series with seasonality

Newsletter – October 2024

The importance of advanced models that can explain and forecast the key financial variables, such as product interest rates, PnLs, volumes, income has been growing significantly in the past years and is expected to so in the future. A special type of these that exhibit complex seasonality, such as daily deposit amounts, transaction volumes is of the interest of liquidity risk and interest rate risk in the banking book (IRRBB) measurement and management. Understanding the underlying process enables the development of more accurate models and creates added value by better decision making. In this newsletter, we will present a practical example using debit and credit card payment data from Ireland. Regarding the broader risk management context, see the newsletter on ICLAAP of September 2021 and IRRBB of April 2024 that are related to these topics.

Typical challenges

Intraday and daily liquidity management at banks

Banks face the challenge of forecasting intraday and daily fluctuations in deposit volumes, cash inflows and outflows, which are crucial for an effective liquidity management. Using advanced models, banks can estimate these movements by capturing recurring patterns such as payday deposits, weekend effects, and seasonal promotions. Effective management of intraday liquidity risk is paramount not only for regulatory compliance but also for the effective utilization of liquidity reserves.

For example, in Hungary, the deposit amounts from natural persons show clearly recognizable seasonal patterns. Just mentioning a few factors, 1) the salaries are usually transferred at the final, or in the beginning days of each month, 2) the pensions are normally transferred on the 12th calendar day of each month[1], and 3) the salary premiums are often transferred in December before Christmas.

Utilization of central bank tools

Central banks offer various tools, such as standing facilities and open market operations, to help banks manage liquidity. Accurate seasonal forecasts enable banks to plan their usage of these tools more effectively, balancing their liquidity needs and regulatory requirements.

For example, the system of the required reserves has a monthly cycle. A more accurate forecast for the major components of the bank's own, and banking sector's liquidity can provide higher quality information for a grounded decision making.

1 For the exact schedule of pension payments in 2024, see the relevant webpage of Hungarian State Treasury (Magyar Államkincstár)

Card Settlement Forecasting

Accurate forecasting of card settlement volumes is essential for financial institutions and payment processors. Banks can model daily and seasonal variations in transaction volumes, driven by factors such as holiday shopping, sales events, and travel seasons.

For example, a payment processor can predict a spike in transaction volumes during Black Friday and the holiday season. By forecasting these trends, they can ensure adequate processing capacity and liquidity, avoiding delays and enhancing customer satisfaction.

Case study

Used Data

As a demonstration for the techniques that can provide to the challenges described above, we analyzed a [dataset on daily expenditure records from credit and debit cards issued to Irish residents]. It includes total daily spending on these cards and ATM withdrawals from March 2020 to August 2022. This time series exhibit many phenomena, such as weekly, monthly and yearly seasonal patterns, special behavior related to Easter and other holidays, and an overall trend that is tied to the general sales turnover, which also affected by COVID. The time series is visually shown in the Figure 1.

We used the [monthly retail total sales index] as an external exploratory variable in the analysis. However, its frequency is lower than the dependent variable, therefore we transformed it to a higher frequency. We note that there are advanced techniques to align time series data of different frequencies, ensuring that the sum, average, first, or last value of the resulting high-frequency series aligns with the original low-frequency data.

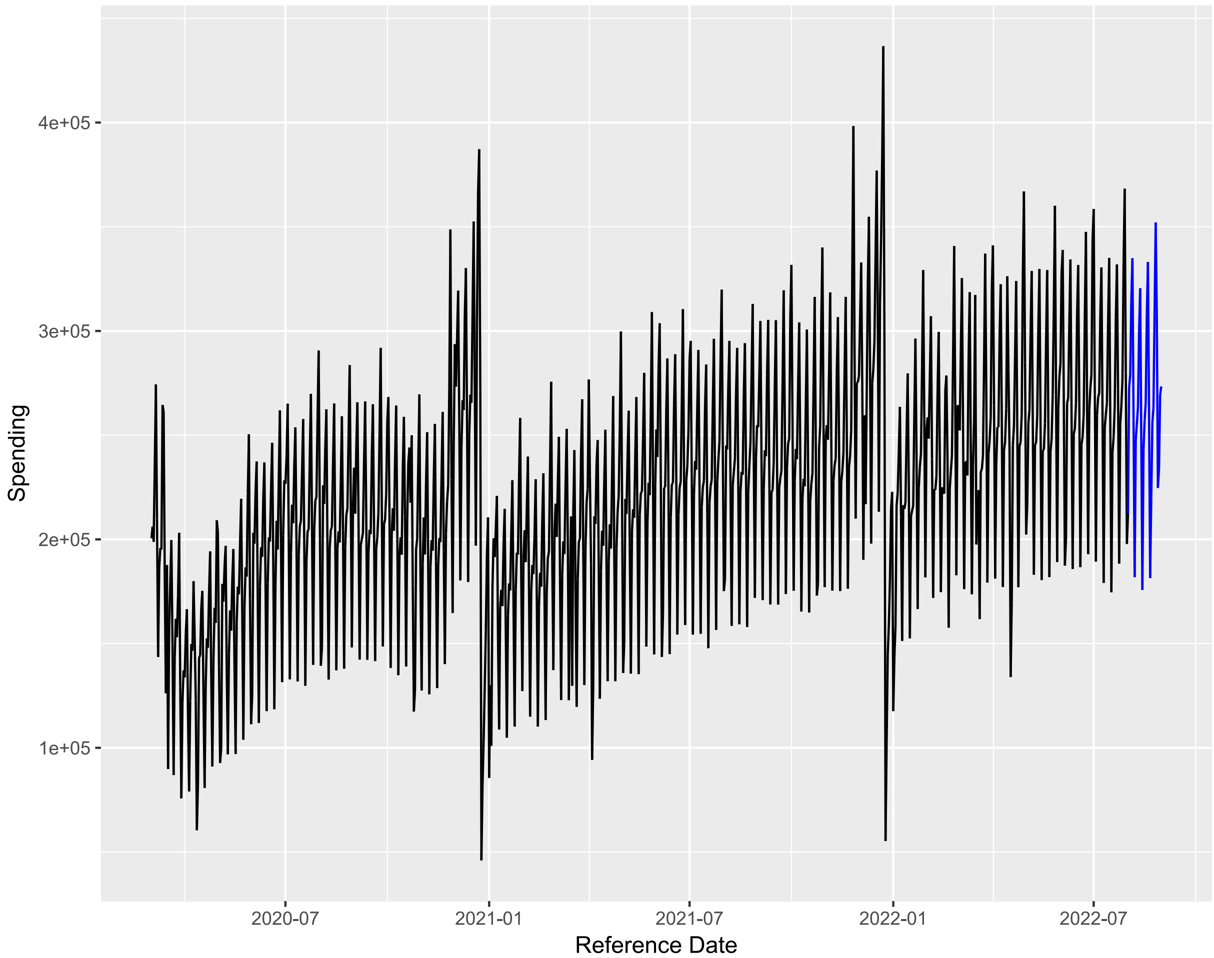

Figure 1.
The time series of debit and credit card spendings, the black section indicates the training data set, and blue section indicates the test data set

Used models

Naïve approaches

The naïve approaches offer a simple, effective rule for generating future predictions and provides a baseline for comparing more complex models. The simple naive model uses the most recent data point for all future predictions but leads to recency bias. In contrast, the seasonal naïve model improves this by using the last observed value from the same period in the past, such as using last week's Friday sales to forecast sales for the upcoming week's Friday.

Seasonal Autoregressive Integrated Moving Average (SARIMA) model

For the more sophisticated models, we used several explanatory variables. We incorporated seasonal variations into our analysis, including weekly, monthly, and yearly patterns with the help of a special family of functions handling seasonality. Additionally, we introduced various dummy variables to account for outliers that occur during specific holidays, such as Christmas, and on specific Irish holidays like March 17th (Saint Patrick's Day). We constructed the trend using multiple linear segments, each representing the trend for the corresponding quarter. Before using the model for forecasting, we ensured the time series is stationary by removing any trend or seasonality through differencing, including seasonal differencing if necessary. Pros of SARIMA include its proficiency in short-term forecasting using solely historical data and its capability to handle non-stationary data effectively. However, SARIMA models are limited by their poor performance in long-term forecasting, difficulty in predicting turning points, computational demands, and subjective parameter selection process.

Time series linear model (TSLM)

These models are used to fit linear models to time series including trend and seasonality components. When we added a lagged component of the data, we achieved even more better results. We also used piecewise linear trend, special family of functions to handle seasonality, and the external variable similarly to the SARIMA model.

Neural network autoregression (NNAR) model

NNAR models utilize lagged values of a time series as inputs within a neural network framework. This approach is analogous to traditional autoregression but incorporates nonlinear processing through hidden layers. An NNAR model can incorporate both non-seasonal and seasonal inputs, with configurations denoted by parameters indicating the number of lagged observations and the number of neurons in the hidden layer.

Prophet model

The Prophet model, developed by Facebook, is designed for forecasting time series data that exhibit strong seasonal patterns and are influenced by holidays. It functions as a nonlinear regression model that incorporates a piecewise-linear trend, seasonal fluctuations, and holiday effects represented by dummy variables. This model is particularly effective for daily data with multiple seasons of history, using a Bayesian framework for flexible model fitting and automatic determination of trend changes and other parameters.

Combination

Forecast combination is a technique where the outputs of multiple models are combined to improve prediction accuracy. By aggregating the forecasts of individual models, combined forecasts can often achieve better results than any single model alone. This phenomenon has been observed by many times. We defined combination as the simple arithmetic mean of the other models we used.

Evaluating the Results

Training and test data sets

We used data from March 2020 to July 2022 as the training dataset, and the remaining month, August 2022 as the test dataset. The training set is used to train our model, while the testing set is used to evaluate its performance. This approach assesses how well our model is likely to perform on unseen data and helps identify any issues, such as overfitting. We note that our analysis is made for demonstrational purposes only, it is not comprehensive. It should not be used to assess the time series or models, or make forecasts in real business context.

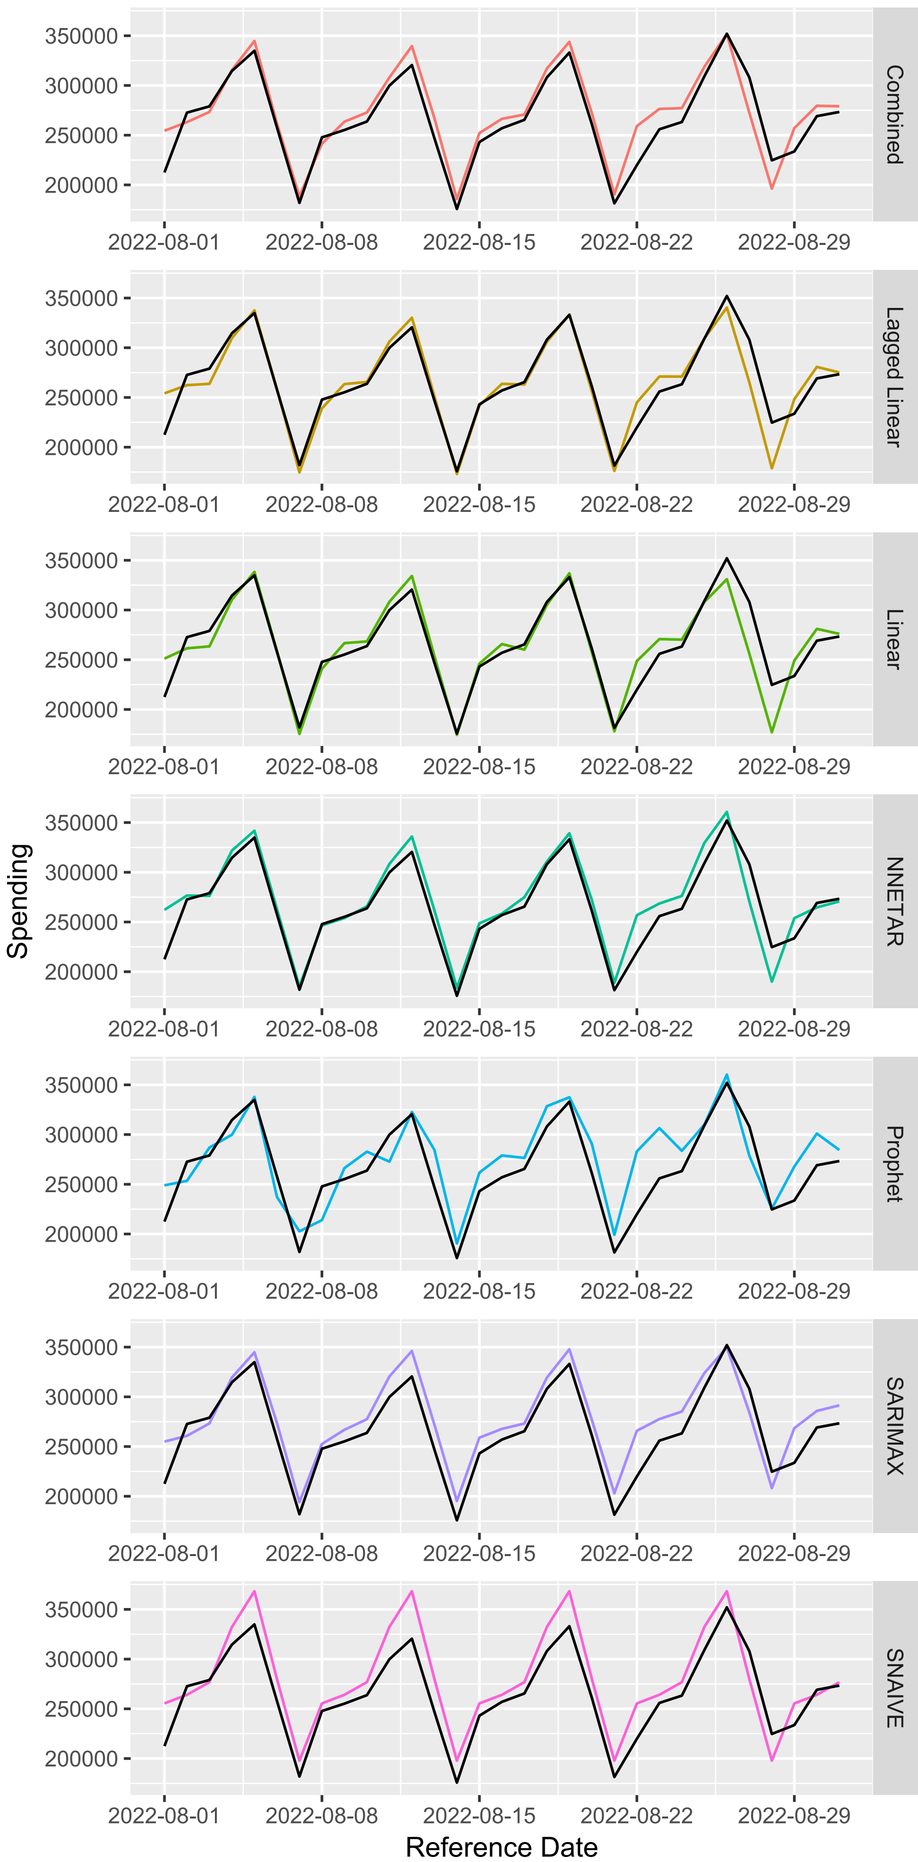

Figure 2

The forecasts in the test data set

Metrics

The following forecast accuracy metrics were applied to evaluate the models by comparing forecasted values to actual observations. Mean Error (ME) assesses model bias, while Root Mean Square Error (RMSE) and Mean Absolute Error (MAE) quantify the magnitude of errors, with RMSE being more sensitive to large errors due to its squaring of residuals. Mean Absolute Percentage Error (MAPE) and Mean Percentage Error (MPE) express errors as a percentage of actual values, making them intuitive for business applications. Additional metrics like Mean Absolute Scaled Error (MASE) and Root Mean Square Scaled Error (RMSSE) provide scale-independent error comparisons, and the First Autocorrelation of Forecast Errors (ACF1) detects patterns in residuals that may indicate model deficiencies.

Results

The results through the performance metrics are shown in Table 1 for the training dataset, and in Table 2 for the test data set. The results are also shown visually in Figure 2. As a benchmark, the performance metrics for the SNAIVE approach were also calculated. This is a basic forecasting approach, in our example, all except the Prophet model were able to perform better than it.

The best performing model in the training data set is the NNETAR based on most of the metrics, however, this model performs only average on the test data and is a typical example of overtraining.

In the test data set, the TSLM (with lag) has the best performance for all the metrics except the 1st order autocorrelation, which suggest that there could unexploited patterns. It is also very important that this is much better than the SNAIVE approach, which is considered to be the baseline forecast here.

Model	Type	ME Mean Error	RMSE Root Mean Square Error	MAE Mean Absolute Error	MPE Mean Percentage Error	MAPE Mean Absolute Percentage Error	ACF1 1st order auto-correlation
Combined	Training	641.46	15381.09	9153.36	-0.61	4.77	0.47
TSLM (with lag)	Training	511.13	14624.24	9439.11	-0.32	4.76	0.07
TSLM (without lag)	Training	676.19	18423.46	11937.52	-0.45	5.90	0.49
NNETAR	Training	400.47	**7809.49**	**4810.78**	**-0.05**	**2.17**	0.06
Prophet	Training	**6.05**	18376.19	11582.86	-0.77	5.97	0.47
SARIMAX	Training	836.13	15823.76	9872.21	-0.37	5.22	**0.00**
SNAIVE	Training	577.51	33870.81	18553.44	-1.89	10.23	0.61

***Table 1:** The performance metrics in the training data set*

Model	Type	ME Mean Error	RMSE Root Mean Square Error	MAE Mean Absolute Error	MPE Mean Percentage Error	MAPE Mean Absolute Percentage Error	ACF1 1st order auto-correlation
Combined	Test	-8319.84	17864.61	13899.89	-3.45	5.54	0.26
TSLM (with lag)	Test	**244.75**	**15989.90**	**10461.09**	**-0.03**	**4.20**	0.28
TSLM (without lag)	Test	-46.44	17514.31	11852.77	-0.20	4.64	0.37
NNETAR	Test	-8491.86	16828.46	12570.68	-3.38	4.94	**0.09**
Prophet	Test	-12870.67	30317.57	23948.79	-5.59	9.58	0.24
SARIMAX	Test	-13468.09	19967.65	17348.88	-5.57	7.00	0.30
SNAIVE	Test	-15286.72	23022.96	19844.68	-5.92	7.65	0.32

***Table 2:** The performance metrics in the test data set*

Further steps needed for real-life setting

Because of length constraints, we conclude our newsletter here, but we note that in practice, a proper model development process would involve additional steps for a more comprehensive analysis and more detailed understanding of the data generating process. For example, the Time Series Cross-Validation (TSCV) is a robust technique for evaluating models in time series analysis, ensuring that the temporal order of observations is preserved. Unlike traditional cross-validation, which uses random data splits, TSCV maintains the temporal order of observations, training on past data and testing on future data to mimic real-world scenarios. This method helps in the assessment whether the model produces stable outputs over time or not.

The newsletter was prepared by: József Soltész, Judit K. Nagy, György Székely.

Contacts:

Ágnes Rakó
Partner
M: +36 70 370 1792
E: agnes.rako@kpmg.hu
KPMG.hu

Péter Szalai
Associate Partner
M: +36 70 370 1739
E: peter.szalai@kpmg.hu

József Soltész
Senior Manager
M: +36 70 370 1766
E: jozsef.soltesz@kpmg.hu

BLOG