9. Demand forecasting, visualising and feature engineering of time series in supply chains

What is demand forecasting? How can we effectively visualize and make conclusions about the customer data? How to conduct the feature engineering of time series?

On these and similar questions, we will try to provide answers in the following chapter.

9.1 What is customer demand and demand forecasting?

The final customer's demand sets the entire supply chain in motion stakeholders (Syntetos et al., 2016). Accordingly, customer demand is a key component for planning all logistic processes in the supply chain, and therefore determining levels of customer demand is of great interest for supply chain managers. Complementary, demand forecasting is an essential activity for planning and scheduling logistic activities within the observed supply chain (Mircetic et al., 2017). Accurate demand forecasting models directly influence the decrease of logistics costs, since they provide an assessment of customer demand (Mircetic et al., 2016). Forecasting in supply chains goes beyond the operational task of extrapolating demand requirements at one echelon. It involves complex issues such as supply chain coordination and the sharing of information between multiple stakeholders (Syntetos et al., 2016).

Customer demand and accompanying forecasts are vital to SCs, as it provides the basic inputs for the planning and control of all functional areas, including logistics, marketing, production, etc (Mircetic, 2018). If the final consumers' demand were constant, or known with certainty well in advance, then the operation of a supply chain would be a straightforward (backwards) scheduling exercise. However, demand is not known and thus it needs to be forecasted. It is the uncertainty associated with this demand that makes supply chain management very difficult (Syntetos et al., 2016). The effectiveness of demand forecasting is influenced by the inherent uncertainties in the demand time series that supply chains have (Rostami-Tabar, 2013). Consequently, addressing and understanding these uncertainties is a major challenge for managers when coordinating and planning operations within supply chains (Mircetic, 2018).

Demand uncertainty is one of the most significant challenges for modern supply chains. The recent COVID-19 pandemic has further underscored this issue, causing widespread disruptions that have complicated supply chain planning and control (Nikolopoulos et al., 2020). Demand forecasting in supply chains often involves predicting the demand for numerous items. Forecasters in supply chains typically extrapolate time series data for each stock-keeping unit individually. For example, a retailer might use point-of-sale data to generate forecasts at the individual store level (Mircetic et al., 2022).

9.2 Demand forecasting steps in supply chains?

In accordance with the statements and conclusions made above, regarding the importance of demand and demand forecasting for supply chains, it is important to follow the specific steps when developing the forecasts in supply chains (Figure 9.1).

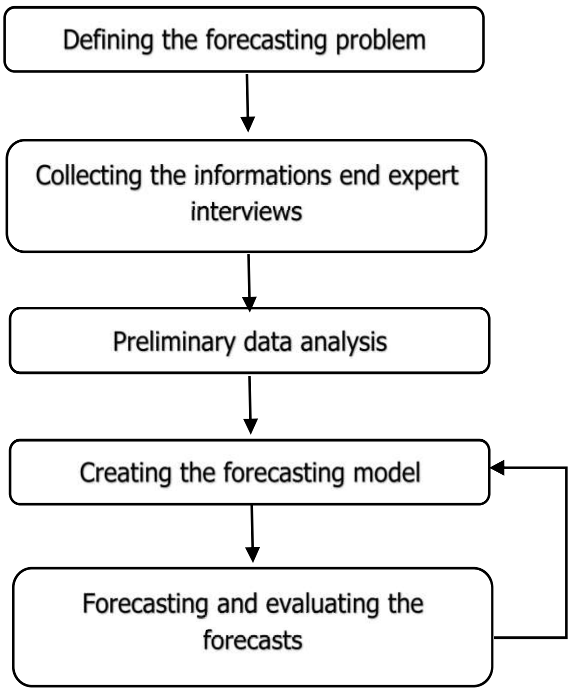

Figure 9.1 Thee basic steps for proper implementation of forecasts within a company (Makridakis et al., 1998; Makridakis et al., 1983).

Each of the steps in Figure 9.1 has its merits and contribution towards creating reliable and useful (business-oriented) forecasts. Accordingly, defining the problem is often the most difficult part of forecasting and requires an understanding of how the forecasts will be used, as well as the role of forecasting functions within the observed company. The forecaster should spend considerable time communicating with everyone involved in data collection, database maintenance, and using forecasts for future planning. One of the main aggravating factors in problem definition is how the final forecast will be utilized in everyday logistics operations (what platform, software design, user interface, etc).

For the information collection step, there are at least two types of information are always needed: statistical data and the accumulated expertise of the people who collect the data and use the forecasts. In practice, it is often difficult to obtain historical data to create a good statistical model. Also, there is a big misunderstanding of what is demand data and what can it be used as its proxy. There is a bad practice of using the shipment and delivery data as a proxy for demand data, which will only deteriorate the decision-making process based on forecasts made on the easy kind of data. Sales data is the only reliable proxy for demand data (Syntetos et al., 2016), although this simplification is not perfect, especially in supply chains with a lot of out-of-stock situations.

For the preliminary data analysis step, it is recommended to always start data analysis with graphical representations to answer the following questions. Are there consistent patterns? Is there a significant trend? Is there noticeable seasonality? Is there evidence of business cycles? How strong are the relationships between variables? These are the questions on which simple graphics can provide answers and allow further data analysis by narrowing the focus of which models to apply to discovered demand features. Usually, the simple model, determined in this way can beat the more sophisticated and complicated ones (Rostami-Tabar & Mircetic, 2023).

Selecting and creating forecasting models is the most important step when creating the forecasting model. Which model to use depends on several factors, the most important of which are the availability of historical data and the correlation between dependent and independent variables. It is common to compare two or three potential models when selecting a model. Each model is an artificial construct based on a set of assumptions (explicit and implicit) and generally involves one or more parameters that must be created using known historical data. In the following chapter, the process of development and application of the ARIMA model will be represented, as one of the best and most popular models.

Evaluating the forecasting model is the step which measures the usability of the created model. After selecting the forecasting model and estimating its parameters, the model is used to create forecasts. The accuracy of the model is evaluated using various statistics, but it is also important to test forecasts through business implication measures (i.e., utility metrics).

9.3 Demand forecasting in the food industry

The consumption data for all products from the observed company in the food industry is presented as weekly demand spanning from January 2012 to December 2014 (Figure 9.2). The x axis represents the time, while the demand values are presented on y axis. This data is shown in weekly intervals, as this corresponds to the period during which the supply to final points of sale is conducted. Consequently, the company's management is focused on forecasting the market's weekly consumption.

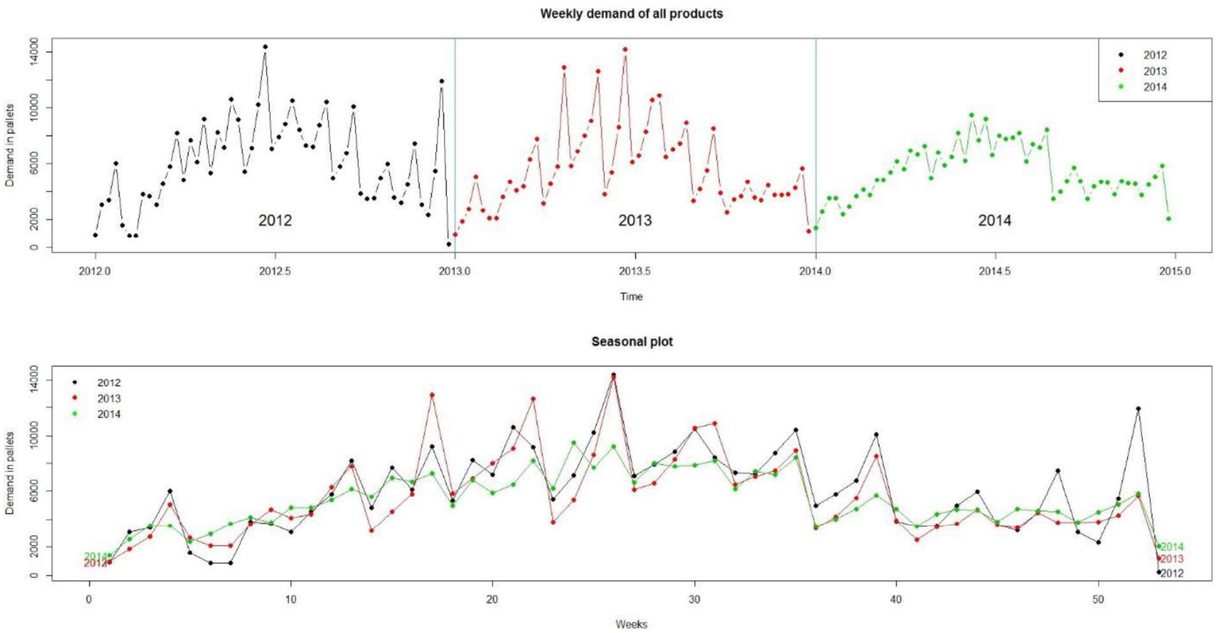

Figure 9.2 The demand data for observed food company.

Figure 9.2 demonstrates several important features and characteristics about the given data. First, the data has a strong seasonal character with peak sales occurring in the middle of the year (summer months). Secondly, the sub-seasonal plot (bottom plot in Figure 9.2) displays a significant pattern change of the downturn trend in 2014! These is very significant characteristics for choosing the right forecasting model and for senior managers in the company since it reveals a significant drop in consumption and market loss. To further investigate the observed characteristics the seasonal trend decomposition (STL) methodology is used (Figure 9.3). The STL decomposition divides the original demand patterns in three components: seasonal, trend and remainder component.

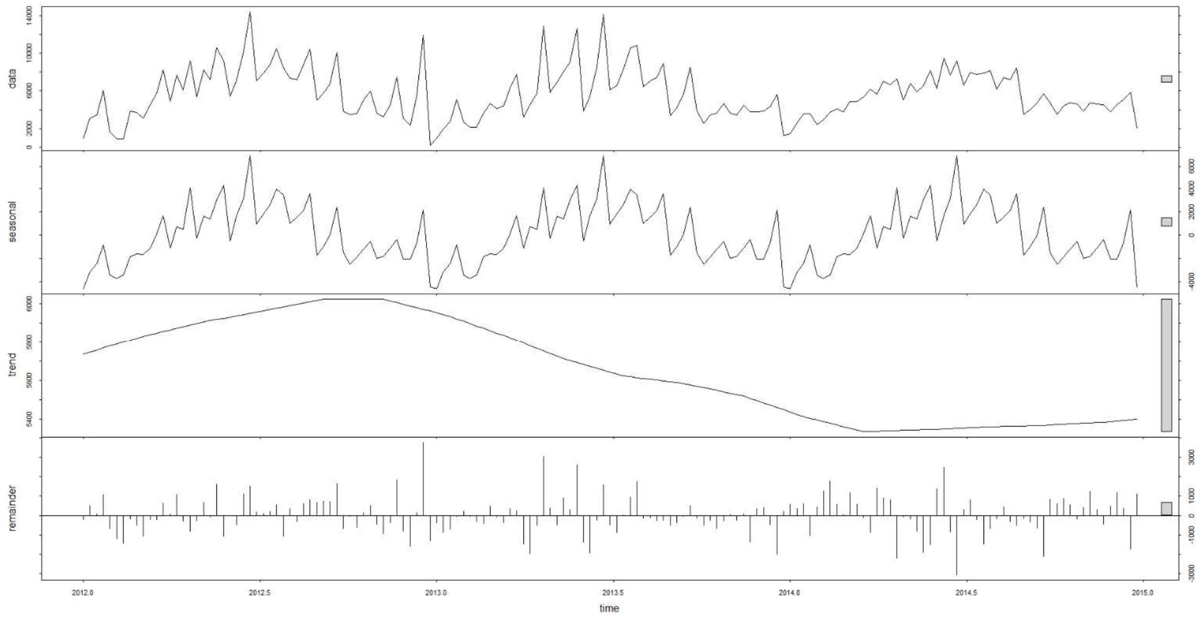

Figure 9.3 The STL decomposition of demand data.

The STL revealed that observed time series exhibit an additive nature, meaning that fluctuations around the trend-cycle curve do not significantly increase over time. As a result, Box-Cox transformations were not necessary for the raw time series. Decomposition revealed that the seasonal component is dominant in the observed series, showing high fluctuations within one year. Given the limited number of years of observations, it is difficult to identify a business cycle. Decomposition also indicated that the trend in the series is minimal, with a decreasing pattern starting from mid-2013.

These identified characteristics represented the important input during the process of design of the appropriate forecasting model. For this purpose, the S-ARIMA (Seasonal Autoregressive Integrated Moving Average) model is chosen. The S-ARIMA model is highly effective for forecasting because it combines both autoregressive and moving average components, along with differencing to make the data stationary. This model is particularly adept at capturing and modelling seasonal patterns in time series data, making it ideal for industries with cyclical demand patterns, such as the food industry. Additionally, S-ARIMA's ability to handle complex seasonal structures and trends allows for more accurate and reliable forecasts, which are crucial for effective supply chain planning and inventory management. The S-ARIMA structural form is presented in Equation (1).

$$\Phi(B^m)\phi(B)(1-B^m)^D(1-B)^d y_t = c + \Theta(B^m)\theta(B)\varepsilon_t, \qquad (1)$$

where $\Phi(z)$ and $\Theta(z)$ are polynomials of order P and Q respectively, each containing no roots inside the unit circle. B is the backshift operator used for describing the process of differencing, i.e. $By_t=y_{t-1}$. If $c \neq 0$, there is an implied polynomial of order $d+D$ in the forecast function. Since the S-ARIMA is a highly parameterized model, the key question when using the S-ARIMA model is selecting the appropriate model order, which involves determining the values of p, q, P, Q, D, and d. If d and D are known, the orders p, q, P, and Q can be selected using an information criterion such as the Akaike Information Criterion (AIC) or Bayesian Information Criterion (BIC). The formulas for AIC and BIC are given by:

$$AIC = -2\log(L) + 2(k);$$

$$BIC = N\log\left(\frac{SSE}{N}\right) + (k+2)\log(N) \qquad (2)$$

where $k=p+q+P+Q+1$ if a constant term is included and 0 otherwise, L is the maximized likelihood of the model fitted to the differenced data, SSE is the sum of squared errors, N is the number of observations used for estimation, and k is the number of predictors in the model.

For determining the optimal set of parameters Hyndman and Khandakar (2007) proposed the Canova-Hansen and KPSS unit root test in the following steps:

- Use the Canova-Hansen test for determining D in the ARIMA framework.
- Choose d by applying a successive KPSS unit root test on seasonally differenced data (if $D = 1$) or on the original data (if $D = 0$).
- Select the optimal values for p, q, P and Q by minimizing the AIC.

9.4 Developing the S-ARIMA forecasting model

Developing the S-ARIMA forecasting model

In accordance with the procedure mentioned above several parameter settings are tested and their performance is presented in Table 9.1. For testing the performances of different models several measures are used: mean absolute percentage error (MAPE), root mean square error (RMSE), mean absolute scaled error (MASE), AIC and BIC (Equations 2 and 3)

$$MAPE = \frac{1}{N}\sum_{1}^{N}|e_i|;$$

$$RMSE = \sqrt{\frac{1}{N}\sum_{1}^{N}e_i};$$

$$MASE = \frac{1}{N}\sum_{1}^{N}|q_j|, \tag{3}$$

where e_j are residuals and $q_j = \dfrac{e_j}{\dfrac{1}{T-m}\sum_{t=m+1}^{T}|y_t - y_{t-m}|}$.

The most popular measures for the managers in the supply chains are RMSE and MAPE since they provide a "sense" of how good the model is performing in real numbers (RMSE) and percentages (MAPE).

Table 9.1 Performance of S-ARIMA models with different parameter settings.

Models[a]	RMSE	MAPE	MASE	AIC	BIC
S-ARIMA $(5,0,1)(1,0,0)_{52}$[b]	1023	14.18 %	0.60 %	86.62	110.50
S-ARIMA $(4,0,0)(1,0,0)_{52}$[c]	1041	14.53 %	0.62 %	86.73	105.31
S-ARIMA $(4,0,0)(0,1,1)_{52}$[d]	1882	22.12 %	1.05 %	41.74	53.560
S-ARIMA $(4,0,1)(1,0,0)_{52}$[e]	1050	14.18 %	0.60 %	88.23	109.47
S-ARIMA $(0,0,1)(0,1,0)_{52}$[f]	1797	21.42 %	1.02 %	36.52	40.460

[a] The model errors are calculated on the test data set.

[b] Detalis about S-ARIMA $(5,0,1)(1,0,0)_{52}$ model are provided bellow.

$_c(1-0.39B+0.06B+0.04B-0.46B)(1-0.65B^{52})y_t = 8.52$.

$_d(1-0.29B+0.19B-0.05B-0.19B)(1-B^{52})y_t = (1+0.12B^{52})e_t$.

$_e(1-0.29B+0.01B+0.05B-0.48B)(1-0.65B^{52})y_t = 8.52+(1+0.13B)e_t$.

$_f(1-B^{52})y_t = (1+0.3B)e_t$.

Table 9.1 demonstrates that the S-ARIMA $(5,0,1)(1,0,0)_{52}$ model outperformed the competing models by achieving the lowest RMSE, MAPE, and MASE errors. The S-ARIMA $(5,0,1)(1,0,0)_{52}$ model form is presented in Equation (4), where $\phi_1 = -0.5421$, $\phi_2 = 0.2962$, $\phi_3 = -0.099$, $\phi_4 = 0.3974$, $\phi_5 = 0.4994$, $\Phi_1 = 0.9558$, $c = 8.523$, and $\Theta_1 = 0.6345$.

$$(1-\phi_1 B-\phi_2 B-\phi_3 B-\phi_4 B-\phi_5 B)(1-\Phi_1 B^{52})y_t = c+(1+\theta_1 B)e_t\,, \tag{4}$$

Similar performance was observed with the S-ARIMA $(4,0,0)(1,0,0)_{52}$ and S-ARIMA $(4,0,1)(1,0,0)_{52}$ models. When performing forecasts the S-ARIMA $(5,0,1)(1,0,0)_{52}$ model on average produces an error of 1023 products, which translates to 14.18 %. These results highlight the importance of carefully selecting the terms to include in an S-ARIMA model since there are no significant differences between the performances of the model with different parameters. The evaluation indicated that models incorporating autoregressive (*p, P*) and moving average components (*q, Q*) were more effective in forecasting beverage consumption than those incorporating seasonal or non-seasonal differencing. Additionally, the comparative review revealed some unexpected findings, such as the models including seasonal differencing S-ARIMA $(4,0,0)(0,1,1)_{52}$ and S-ARIMA $(0,0,1)(0,1,0)_{52}$ performing worse than a simple average naive forecast model, as evidenced by their MASE errors exceeding one.

9.5 Forecasts of the future demand

Figure 9.4, demonstrates the performance of S-ARIMA $(5,0,1)(1,0,0)_{52}$. The top left panel represents the starting demand data, coloured for easier distinction of different marketing years. The top right panel is the input data to the S-ARIMA $(5,0,1)(1,0,0)_{52}$, representing the training data on which parameters in S-ARIMA are determined following the procedure described in subchapter 9.3. This is very important to understand how difficult the job of the forecasting model is, since in this case it has two years of data as input and needs to forecast future demand one year ahead! This is quite a common case scenario in supply chains and logistics since there is an unwritten rule that companies keep a history of their data for three years after which they discard the data.

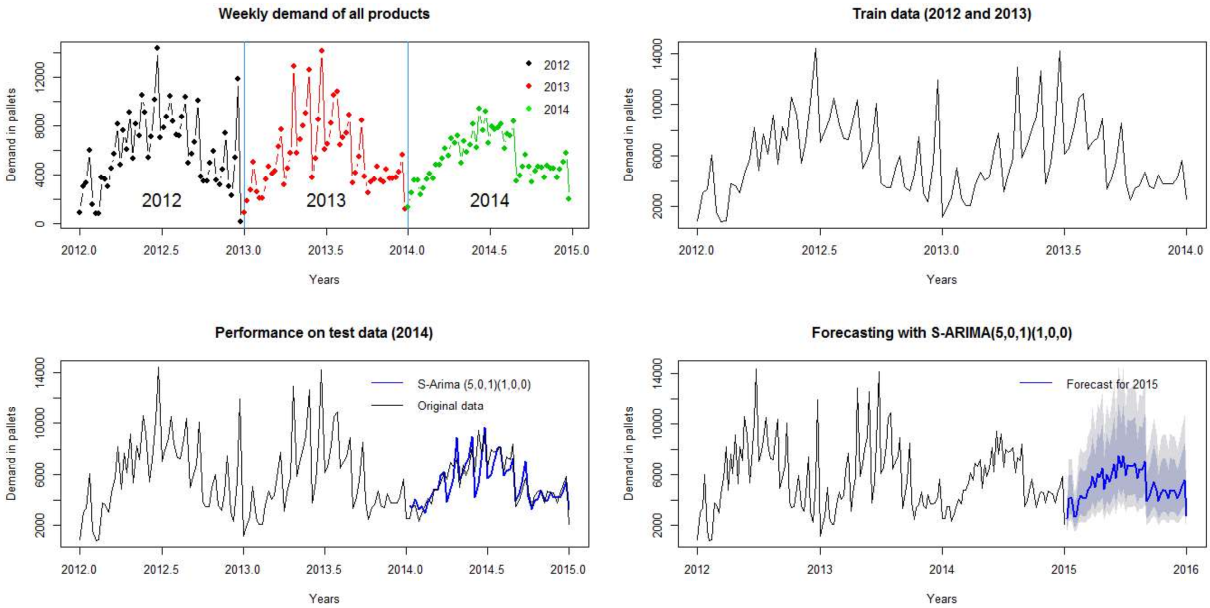

Figure 9.4 Train, test and forecasted demand of S-ARIMA (5,0,1)(1,0,0)$_{52}$ model.

The bottom left panel in Figure 9.4, presents the performance of an observed model. Model performance is already presented in Table 9.1 via different statistical measures, but for managers is usually hard to get the feeling of how good or bad the model is. For this purpose, the panel graphically demonstrates the performance of the model. It could be argued that the model follows the test data pretty well and in the majority of periods demonstrates excellent performance. In order to generate future forecasts, S-ARIMA (5,0,1)(1,0,0)$_{52}$ refitted by adding the data from 2014 to the train data. After that model produced 52- weekly steps ahead forecasts for 2015. The forecasts for 2015 are presented in the bottom left panel. The forecasts are accompanied by 80% and 95% prediction intervals, which demonstrate possible future forecasts disperse from the mean predicted values. The model forecasts a continued decline in demand that began in 2014. The causes of this decline could be diverse and should be further explored by managers at the strategic level of the company.

References Chapter 9

- Hyndman, R., & Khandakar, Y. (2007). Automatic time series forecasting: The forecast package for R. *Journal of Statistical Software, 27*(3). http://www.jstatsoft.org/v27/i03
- Makridakis, S., Wheelwright, S., & Hyndman, R. J. (1998). *Forecasting: Methods and applications* (3rd ed.). New York: John Wiley & Sons.
- Makridakis, S., Wheelwright, S., & McGee, V. (1983). *Forecasting: Methods and applications* (2nd ed.). New York: John Wiley & Sons.
- Mircetic, D. (2018). *Boosting the performance of top-down methodology for forecasting in supply chains via a new approach for determining disaggregating proportions*. University of Novi Sad.
- Mircetic, D., Nikolicic, S., Maslaric, M., Ralevic, N., & Debelic, B. (2016). Development of S-ARIMA model for forecasting demand in a beverage supply chain. *Open Engineering, 6*(1).
- Mircetic, D., Rostami-Tabar, B., Nikolicic, S., & Maslaric, M. (2022). Forecasting hierarchical time series in supply chains: An empirical investigation. *International Journal of Production Research, 60*(8), 2514-2533.
- Mircetic, D., Nikolicic, S., Stojanovic, D., & Maslaric, M. (2017). Modified top-down approach for hierarchical forecasting in a beverage supply chain. *Transportation Research Procedia, 22*, 193–202.
- Nikolopoulos, K., Punia, S., Schäfers, A., Tsinopoulos, C., & Vasilakis, C. (2020). Forecasting and planning during a pandemic: COVID-19 growth rates, supply chain disruptions, and governmental decisions. *European Journal of Operational Research, 290*(1), 99–115.
- Rostami-Tabar, B. (2013). *ARIMA demand forecasting by aggregation*. Université Sciences et Technologies Bordeaux I.
- Rostami-Tabar, B., & Mircetic, D. (2023). Exploring the association between time series features and forecasting by temporal aggregation using machine learning. *Neurocomputing, 548*, 126376.
- Syntetos, A. A., Babai, Z., Boylan, J. E., Kolassa, S., & Nikolopoulos, K. (2016). Supply chain forecasting: Theory, practice, their gap and the future. *European Journal of Operational Research, 252*(1), 1-26.