

Booking Horizon Forecasting with Dynamic Updating: A Case Study on Hotel Reservation Data

Alwin Haensel

*Department of Mathematics, VU University Amsterdam, ahaensel@few.vu.nl
Bookit B.V., Amstelveen, ahaensel@bookit.nl*

Ger Koole

Department of Mathematics, VU University Amsterdam, koole@few.vu.nl

Abstract

A highly accurate demand forecast is fundamental to the success of every revenue management model. As often required in practice and theory, we aim to forecast the accumulated booking curve as well as the number of expected reservations for each day in the booking horizon. To reduce the high dimensionality of this problem, we apply singular value decomposition on the historical booking profiles. The forecast of the remaining part of the booking horizon is dynamically adjusted to the earlier observations using the penalized least squares and the historical proportion method. Our proposed updating procedure considers the correlation and dynamics of bookings within the booking horizon and between successive product instances. The approach is tested on real hotel reservation data and shows a significant improvement in forecast accuracy.

Key words: Demand forecasting, dynamic forecast updating, dimension reduction, penalized least squares, time series, revenue management.

1. Introduction

Revenue management (RM) methods are used to control the booking process in many industrial sectors, especially in the airline, hotel and car rental industry. A brief introduction to models and topics of RM is given in Talluri and Van Ryzin (2004b) and a detailed research review is presented in Chiang et al. (2007). Hotel specific RM techniques are illustrated in Vinod (2004) with a description of steps and challenges to undertake to build a successful RM system. At the heart of every revenue management model always lies a demand forecast, whose accuracy is crucial for the success of the model. Poelt (1998) estimates for the airline

Preprint submitted to International Journal of Forecasting *October 11, 2010*

industry that a 20% increase in forecast accuracy can be translated in a 1% increase in revenue generated by the underlying RM system. Van Ryzin (2005) and Zeni (2007) argue that new models of demand forecasts are needed to adjust to the new market situation with more competition and less restricted products. Let us refer for example to the airline sector, where new ‘low-cost carriers’ compete with the major airlines by offering unrestricted products at low prices. Traditional RM and demand forecasts are based on a strict segmentation of the demand into disjoint groups. Customers in less competitive markets with highly restricted products are less volatile and even the independency assumption between demand groups can be justified practically by the customer behavior in these market settings. In current times the old assumptions do not fit the market situation any longer and research is needed to adjust the forecasting models and procedure.

Besides the choice of the forecasting model and its adjustment to the demand time series, there are three important steps to include into the forecasting process. The first step is data unconstraining. It is important to note that sales figures are usually not equal to the real demand. This follows from capacity restrictions, booking control actions and the existence of competitors. Second, the customers’ choice behavior has to be considered. A variety of product offers from a company or its competitors influence the customers’ purchasing decision and thus the demand. Different approaches to model the customer’ choice behavior are presented in Talluri and Van Ryzin (2004a) and Haensel and Koole (2010). The third point and the one on which this paper focuses is the dynamic updating of forecasts when new information becomes available. As shown in O’Connor et al. (2000) the forecast accuracy can be improved by updating, especially when the time series is trended. In case of travel, accommodation and holiday products, the usual long booking horizons (plus the dependency of customer decisions due to holidays, special events or super offers) give additional hope for benefits from forecast updating.

Intensive research on forecast updating is done in the context of call centers. A significant correlation between within-day (morning, midday, evening) call arrivals is found. Models are proposed to forecast not only the call volume for future days, but also the updating of expected call volumes for future time periods within the day. In Weinberg et al. (2007) a multiplicative Gaussian time series model, with a Bayesian Markov chain Monte Carlo algorithm for parameter estimation and forecasting, is proposed. Shen and Huang (2008)

suggest a competitive updating method which requires less computation time. Their method consists of a dimensionality reduction of the forecasting problem and a penalized least square procedure to adjust the time series forecast to observed realizations.

In this paper we are adapting the ideas of Shen and Huang (2008) for call center forecasting to the RM context of hotel reservation forecasting. The equivalence to the within-day periods for which the forecast is updated is the booking horizon in the RM setting. In contrast to the call center case, booking horizons for different product instances are overlapping and correlated in their booking pattern and behavior. Another important difference is the level of forecast data. The call volume in call centers is generally very large, compared to often small demand numbers in the revenue management case. In RM problems a forecast on disaggregated level is required, since booking control actions are applied daily and on product level. A detailed description of the hotel reservation forecasting problem and its characteristics with a comparison of basic forecasting methods is given in Weatherford and Kimes (2003). A more advanced model is presented by Rajopadhye et al. (2001), in which they propose long-term forecasting of total reservations by the Holt-Winters method with a combination of booking curve pickup methods for short-term forecasts. More recently, Zakhary et al. (2009) presented a probabilistic hotel arrival forecasting method based on Monte Carlo simulation. All approaches aim to forecast the final reservation numbers, rather than the booking process, which is the focus of this paper.

We were able to work with real hotel reservation data, provided by Bookit B.V., a European short break specialist for hotels and holiday parks with a market leading position in the Netherlands. Since 1989, Bookit has been operating as an independent reservation company for the leisure and business-to-business market.

This paper is organized as follows. First, in Section 2, we introduce and analyze the data. Next, in Section 3, the forecasting methods are explained, followed by the introduction of the forecast updating procedure and methods in Section 4. Finally, in Section 5, numerical results are presented before we conclude our findings in Section 6.

2. Data

For our forecasting analysis we are able to work with real sales data, extracted from three regions A,B and C. The selected regions have very different characteristics such as reservation

volume, city or countryside location and distance from major market. A hotel product is a combination of region, arrival day of week (DOW) and length of stay. For the analysis we consider a booking horizon of four weeks prior to arrival at the hotel, thus the 28th day coincides with the arrival day. The datasets consist of all reservations made for a particular hotel product gathered over 3 years and multiple comparable hotels over all regions. The hotel products are not divided into different price classes, since hotels are interchangeable and products are not distinguished by specific hotels, but by location and hotel standard. To better illustrate the method, we will restrict this analysis to a fixed arrival DOW and length of stay combination. This separation of the forecasting problem into DOWs is widely common in practice, since the reservation patterns and volumes vary significantly for different arrival DOWs, compare for a discussion with Weatherford and Kimes (2003). In research and practice it is common to work on accumulated reservations, i.e., booking curves, rather than on individual reservations per booking horizon day. However, we see two reasons to prefer the latter. First, as a large reservations agency, one rarely runs out of stock. So the primal goal is to maximize the daily number of reservations. Therefore the second visualization form gives a more usable view as to which product to promote. Second, as also stated in Van Ryzin (2005), the current major direction in revenue management research is to incorporate customer choice behavior under offered alternatives. Thus, it is more important to know the expected customer group demand per individual booking day rather than the aggregated totals. We will work with both visualizations of the booking process, using “Acc” and “Ind” to abbreviate the accumulated and individual reservations respectively. Hence, we obtain six datasets on the three regions: A-Acc, A-Ind, B-Acc, B-Ind, C-Acc and C-Ind. All datasets are given in form of a $n \times m$ reservation data matrix X , with $n = 155$ product instances (as rows) and their associated $m = 28$ booking horizon days (as columns). In our case, the product instances correspond to the successive arrival weeks of our hotel products, fixed DOW and length of stay. For clarity, the $X_{i,j}$ entry denotes the number of bookings made for product instance i (arrival week i) at the j^{th} day in the booking horizon. The first 130 product instances/rows are used for data analysis, testing and parameter estimation. The last 22 instances/rows are used in Section 5 for evaluation of the proposed forecast updating methods. There is a gap of four weeks between the estimation and evaluation sample, caused by the time structure in the dataset: At arrival of product i , the realization of the first booking horizon

week of product instance $i + 3$ is known. The booking behavior of the first three instances in A-Acc and A-Ind, i.e., rows of X , are shown in Figure 1. The total aggregated numbers




Figure 1: Example plot of first three product instances / first three rows of A-Acc and A-Ind.

of reservations received for the first 130 product instances of all regions are shown in Figure 2. Note that the time between the product instances, one week, is much smaller than the




Figure 2: Number of total aggregated reservations for the first 130 product instances of all three regions.

booking horizon of four weeks. The mean and variance of the booking behaviors within the booking horizon for all datasets are shown in Figure 2. We observe that the variance is not constant (heteroscedasticity) and that the variance is greater than the mean (overdispersion). In order to stabilize and reduce the variance, we will work on the logarithmic transformed data. Let x denote the number of reservations. Set $y = \log(x + 1)$, one is added because the

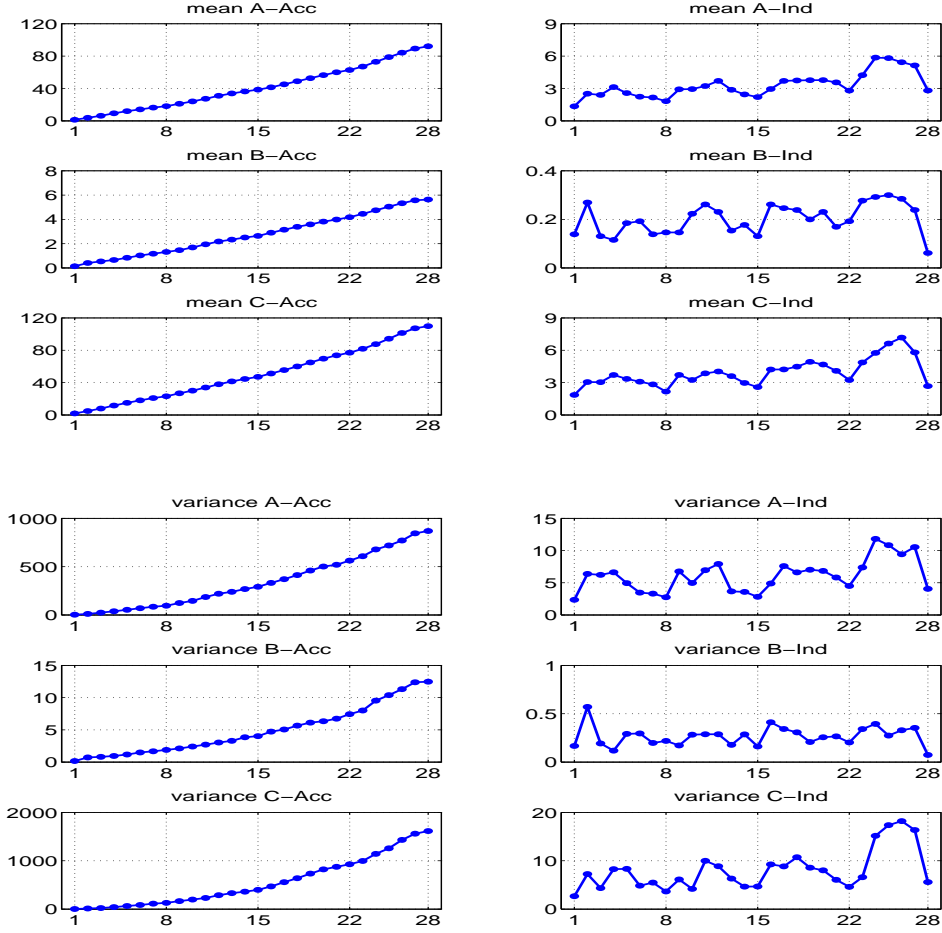


Figure 3: Mean and variance of booking process of all datasets.

dataset contains many booking days with zero reservations. The forecast of y is denoted by \hat{y} and the forecast of x is then given by $\hat{x} = \exp(\hat{y}) - 1$. The following forecasting methods are working on the transformed data, and the forecast error analysis in Section 5 is made on the back transformed data.

An important property in the data structure is the shifted realization time, which means that parts of different product instances are realized at the same time. For example, suppose we select a product instance i , i.e., the i^{th} row in X , and consider the corresponding time as the current time. All information up to instance i plus the first three weeks of the booking horizon of the following instance $i + 1$, the first two weeks of $i + 2$ and the first week of instance $i + 3$ are known at our current time. In other words, fixing an arrival week in our data set as a time point enables us to know the demand realization for the first three weeks of

the booking horizon for the next-week-arrival product. The same is true for the in-two-weeks arrival and in-three-weeks arrival products, where we know the realization of the first two and first week of the booking horizon, respectively. This paper is concerned with the question of how to update reservation forecasts when the realizations of earlier time stages in the booking horizon become known. Therefore we analyze the correlation between reservations made at different moments of the booking horizon.

In Figure 4, the correlation coefficients between early (long in advance) and late reservations (close to arrival at hotel) are plotted as a function of the day in the booking horizon that is the frontier between early and late. The correlation function for split day k is defined on the accumulated reservation dataset X by

$$C(k) = \text{corr}(X_{\cdot,k}, (X_{\cdot,28} - X_{\cdot,k})) \quad k = 1, \dots, 27, \quad (1)$$

where $\text{corr}(a, b)$ is a function returning the linear correlation coefficient between the vectors a and b . The correlation is found to be very different for all three regions. The correlation is highest for region C and lowest for region B. Consequently, the benefit of dynamic forecast updating is assumed to be most beneficial for the datasets of region C. Also the shape of the correlation function differs between the regions. The correlation in region A decreases slightly over the booking horizon, in contrast to regions B and C where the maximum is attained around day 15 (half of the considered booking horizon).

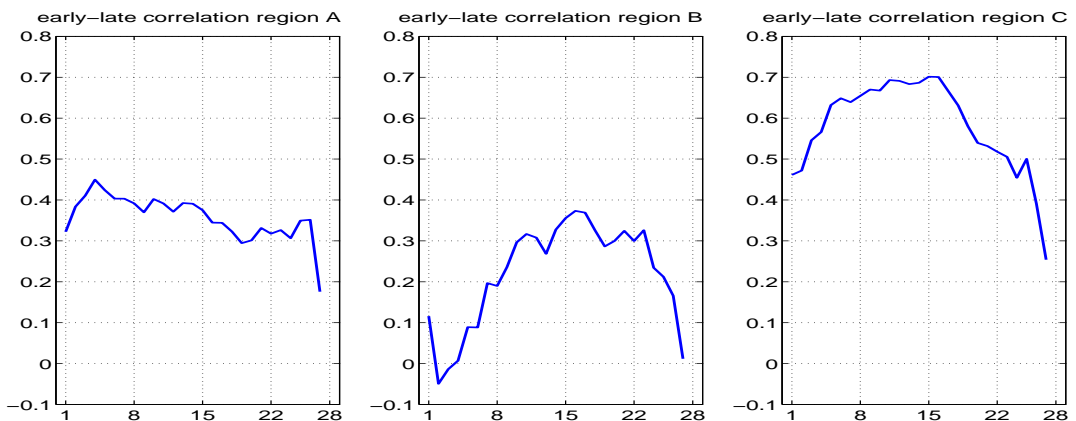


Figure 4: Correlation coefficients $C(k)$ for early-late booking split days $k = 1, \dots, 27$.

Now consider the correlation between bookings in different weeks. Define $w_i = \{7(i-1) +$

$1, \dots, 7(i-1) + 7\}$, the set of days in week i . The correlation function defined on booking weeks w_i and w_j and a dataset X consisting of individual reservations per booking day, is given by

$$C(w_i, w_j) = \text{corr} \left(\sum_{d \in w_i} X_{\cdot, d}, \sum_{d \in w_j} X_{\cdot, d} \right). \quad (2)$$

The correlations are shown in Table 1. Multiple subscripts represent the aggregation over multiple weeks, e.g., $w_{1,2,3}$ stands for the aggregated reservations made in week 1, 2 and 3. It illustrates again the dependence between early and late bookings.

Region	$C(w_1, w_2)$	$C(w_1, w_{2,3,4})$	$C(w_2, w_3)$	$C(w_2, w_{3,4})$	$C(w_3, w_4)$	$C(w_{1,2,3}, w_4)$
A	0.59	0.41	0.59	0.44	0.42	0.27
B	0.00	0.16	0.10	0.24	0.20	0.33
C	0.65	0.64	0.67	0.68	0.46	0.53

Table 1: Correlation $C(\cdot, \cdot)$ of aggregated reservations between specific booking weeks.

3. Forecasting Method

The reservation matrix X contains all reservations per product instance and day in the booking horizon. Our aim is to forecast the number of future reservations to expect for the forthcoming product instances in the next four weeks. For a company that must choose each day which products to offer, advertise or promote, it is very important to know the expected number of reservations per day in the booking horizon. Therefore the forecast is chosen to work on individual numbers of reservations per booking day (“Ind” datasets) as well as on the accumulated reservations (“Acc” datasets). Each of the next four product instances (arrival weeks) has a booking horizon of $m = 28$ days. The i^{th} row of X , $x_i = (x_{i,1}, \dots, x_{i,m})^T$, represents all reservations per day in the booking horizon of instance i . As in Shen and Huang (2008) we are applying singular value decomposition (SVD) to reduce the forecasting dimension. The procedure works as follows:

We are interested in computing a small number of base vectors f_1, \dots, f_K with which the time series $\{x_i\}$ can be reasonably well approximated. The decomposition is given by

$$x_i = \gamma_{i,1}f_1 + \dots + \gamma_{i,K}f_K + \epsilon_i \quad i = 1, \dots, n, \quad (3)$$

where $\gamma \in \mathbb{R}^{n \times K}$ is the weight matrix, $f_1, \dots, f_K \in \mathbb{R}^m$ are the base vectors and $\epsilon_1, \dots, \epsilon_n \in \mathbb{R}^m$ are the error terms. We suppose that the x_i 's can be well approximated by a linear

approximation of the base vectors, so that the error terms are reasonably small. This leads to the following optimization problem

$$\min_{\substack{\gamma_{1,1}, \dots, \gamma_{n,K} \\ f_1, \dots, f_K}} \sum_{i=1}^n \|\epsilon_i\|, \quad (4)$$

for a fixed value K . This problem can be solved by applying SVD to matrix X as follows. Matrix X can be rewritten as

$$X = USV^T, \quad (5)$$

where S is a $m \times m$ diagonal matrix, U and V are orthogonal matrices with dimension $n \times m$ and $m \times m$ respectively. The diagonal elements of S are in decreasing order and nonnegative, $s_1 \geq \dots \geq s_r > 0$, with $r = \text{rank}(X)$ and $s_k = 0$ for all $r + 1 \leq k \leq m$. From (5) we follow now

$$x_i = s_1 u_{i,1} v_1 + \dots + s_r u_{i,r} v_r, \quad (6)$$

where v_k denotes the k^{th} column of matrix V . The K -dimensional approximation is obtained by keeping the largest K singular values ($K < r$), since S is ordered decreasingly the largest are equivalent with the first K values,

$$x_i \approx s_1 u_{i,1} v_1 + \dots + s_K u_{i,K} v_K. \quad (7)$$

Setting now $\gamma_{i,k} := s_k u_{i,k}$ and $f_k := v_k$, for all $i = 1, \dots, n$ and $k = 1, \dots, K$, we have found an optimal solution of (4). The mean squared estimation error (MSEE) of product instance i and fixed K is computed by

$$\text{MSEE}_i = \frac{1}{m} \sum_{j=1}^m \left(x_{i,j} - \left(\sum_{k=1}^K \gamma_{i,k} f_k \right)_j \right)^2. \quad (8)$$

Figure 5 shows the empirical distribution function of the MSEE, computed over the first 130 product instances, for different values of K . We find reasonably small errors for $K = 3$. These values are still outperformed by $K = 5$ or 7 , but for computational reasons we try to keep the dimension small. In the numerical results, where we use $K = 3$ and 5 , we will see that $K = 3$ will produce reasonably good forecasting results. The resulting three base vectors in the case of $K = 3$ and their weights computed over the first 130 instances of the datasets A-Acc and A-Ind are shown in Figure 6. The base vectors represent the data characteristics in decreasing

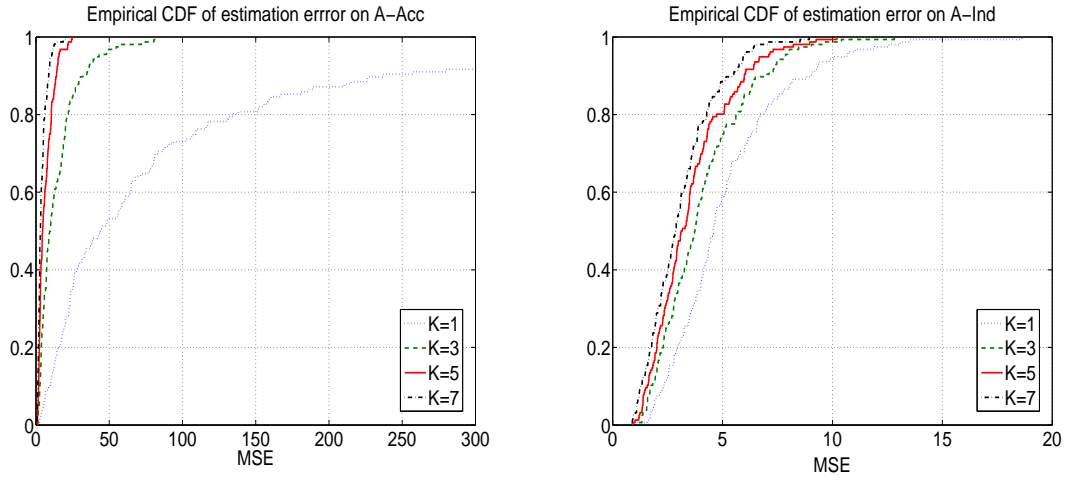


Figure 5: Empirical distribution function of mean squared estimation errors for different numbers of base vectors K .

importance, i.e., the first base vector in A-Ind represents the strong weekly pattern and the first base vector in A-Acc represent the general increasing booking curve pattern. In fact, base vector f_1 in A-Acc is negative and decreasing, but since the corresponding weights time series γ_1 takes negative values, the represented booking pattern is increasing. Remember that the singular value decomposition is applied to the transformed data, when comparing with Figure 2. The forecasting method will work on the time series of $\gamma_{i,k}$ values. The base vectors

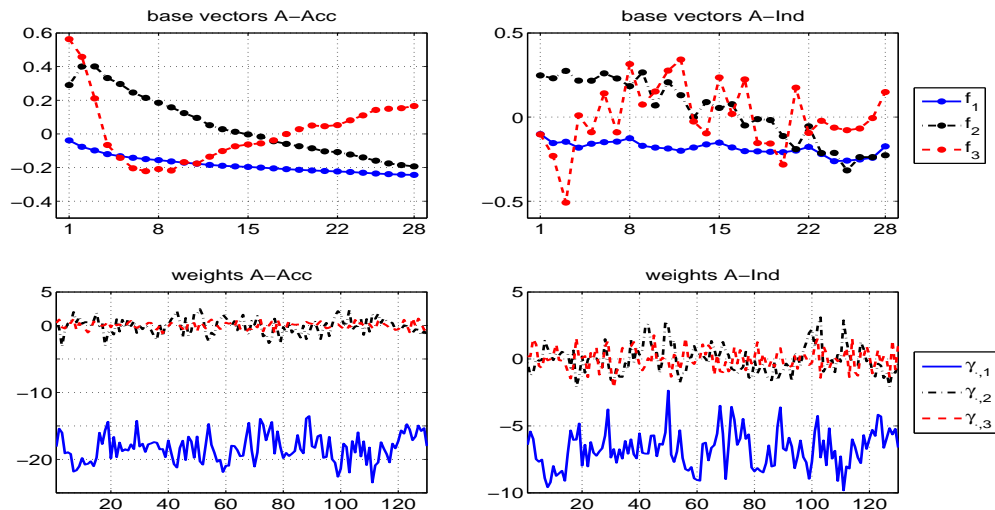


Figure 6: Base vectors f_k and weights $\gamma_{i,k}$ for $K = 3$.

f_1, \dots, f_K are calculated on the historical data and are kept fixed during the forecasting process. Due to the construction of the weights series $\gamma_1, \dots, \gamma_K$ out of the columns of U , we

have that vectors γ_k and γ_l are orthogonal for $k \neq l$. Hence the cross correlation between different weight series can be assumed to be small.

We initially choose as a forecasting method the univariate exponential smoothing with trend and seasonality, i.e., the Holt-Winters (HW) method developed by Holt (1957) and Winters (1960). Holt-Winters is a commonly used method in similar problems, and in practice known to be reasonably accurate, robust and easy to implement. Two seasonal models, additive and multiplicative, are distinguished and both are tested on the three datasets. The additive Holt-Winters (AHW) h -step ahead forecast of $\gamma_{i+h,k}$, for fixed $k = 1, \dots, K$ is

$$\hat{\gamma}_{i+h,k}^{AHW} = a(i) + h \cdot b(i) + c((i+h) \bmod p), \quad (9)$$

where $a(i)$, $b(i)$ and $c(i)$ are given by

$$\begin{aligned} a(i) &= \alpha \cdot (\gamma_{i,k} - c(i-p)) + (1-\alpha) \cdot (a(i-1) + b(i-1)), \\ b(i) &= \beta \cdot (a(i) - a(i-1)) + (1-\beta) \cdot b(i-1), \\ c(i) &= \delta \cdot (\gamma_{i,k} - a(i)) + (1-\delta) \cdot c(i-p). \end{aligned}$$

In contrast, the multiplicative Holt-Winters (MHW) h -step ahead forecast of $\gamma_{i+h,k}$ is

$$\hat{\gamma}_{i+h,k}^{MHW} = \left(a(i) + h \cdot b(i) \right) \cdot c((i+h) \bmod p), \quad (10)$$

where $a(i)$, $b(i)$ and $c(i)$ are computed by

$$\begin{aligned} a(i) &= \alpha \cdot \frac{\gamma_{i,k}}{c(i-p)} + (1-\alpha) \cdot (a(i-1) + b(i-1)), \\ b(i) &= \beta \cdot (a(i) - a(i-1)) + (1-\beta) \cdot b(i-1), \\ c(i) &= \delta \cdot \frac{\gamma_{i,k}}{a(i)} + (1-\delta) \cdot c(i-p). \end{aligned}$$

The period length is one year and because the product instances are weekly $p = 52$. The initial values of a , b and c are derived from a simple decomposition in trend and seasonal component using moving averages (averaging for each time unit over all periods). The decomposition is performed by the R function `Decompose` from the R-stats library. Optimal α , β and δ values are found by minimizing the squared one-step prediction error, evaluated over historical values. Since the weights time series γ_k take negative and positive values, a positive constant is added in the MHW calculation to ensure positivity and subtracted from the forecasts before being processed further. The Holt-Winters forecast for both seasonal

models, of the future booking horizon $\hat{x}_{i+h}^{HW} = (\hat{x}_{i+h,1}^{HW}, \dots, \hat{x}_{i+h,m}^{HW})$ is computed by

$$\hat{x}_{i+h}^{HW} = \hat{\gamma}_{i+h,1}^{HW} \cdot f_1 + \dots + \hat{\gamma}_{i+h,K}^{HW} \cdot f_K. \quad (11)$$

The forecast accuracy for both seasonal models is tested on the sample of the 1-15 step/weeks ahead forecasts, starting at instance 110 (within the estimation sample) and computed on all datasets. The mean squared errors between the actual γ_k and forecasted $\hat{\gamma}_k^{HW}$ are shown in Table 2 and abbreviated with ϵ_k , for $k=1,2,3$. No seasonal model outperforms the other and

Dataset	ϵ_1 AHW	ϵ_1 MHW	ϵ_2 AHW	ϵ_2 MHW	ϵ_3 AHW	ϵ_3 MHW
A-Ind	2.28	2.30	1.23	1.23	1.12	1.12
A-Acc	5.18	5.18	0.79	0.79	0.80	0.80
B-Ind	0.41	0.40	0.29	0.29	0.51	0.51
B-Acc	18.82	19.06	1.42	1.42	2.34	2.78
C-Ind	7.03	7.05	1.87	1.87	0.34	0.40
C-Acc	9.93	10.81	0.90	0.90	0.11	0.11

Table 2: Holt-Winters mean squared errors between the actual and forecasted $\gamma_{\cdot,k}$.

both models produce generally the same forecast error. We will further continue only with the additive seasonal model for the Holt-Winters forecasting method.

Our second forecasting approach is to decompose the γ time series into seasonal, trend and remainder components and to apply an auto-regressive (AR) time series model on the remainder. The additive seasonal models seem to give a good approximation. Therefore we apply a decomposition procedure based on LOESS, i.e., local polynomial regression fitting, as described by Cleveland et al. (1990). The decomposition is performed by the R function STL (seasonal decomposition of time series by loess) from the R-stats library. The γ time series are separately decomposed into additive seasonal, trend and remainder components, see Figure 7 for the case of dataset A-Acc and $K = 3$. The decomposition equation is

$$\gamma_{i,k} = s_{i,k} + d_{i,k} + r_{i,k} \quad \text{for all } i = 1, \dots, n \text{ and } k = 1, \dots, K, \quad (12)$$

where s denotes the seasonal, d the trend and r the remainder component of the time series γ .

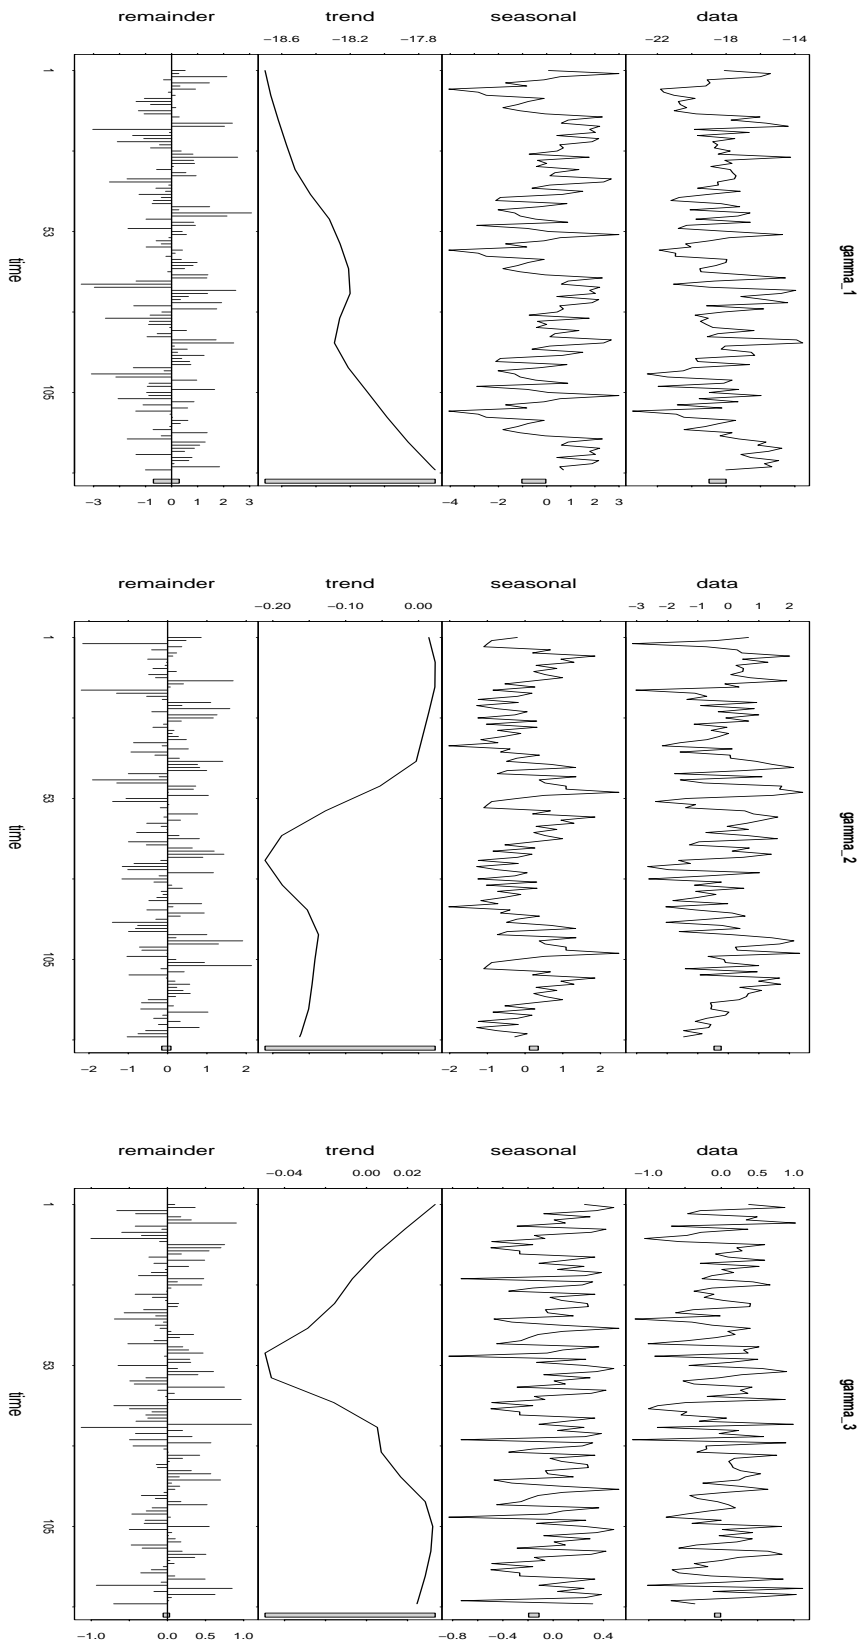


Figure 7: Decomposition of γ time series of the dataset A-Acc into seasonal, trend and remainder for $K = 3$.

The auto-regressive model of $r_{\cdot,k}$ of order p is given by

$$r_{i,k} = \nu_k + a_1 \cdot r_{i-1,k} + \dots + a_p \cdot r_{i-p,k} + u_{i,k}, \quad k = 1, \dots, K, \quad (13)$$

where $a_1, \dots, a_p \in \mathbb{R}$ represents fixed coefficients, u_k a zero mean white noise process and ν_k the intercept. When μ_k denotes the mean of $r_{\cdot,k}$, the intercept is defined as $\nu_k = (1 - \sum_{t=1}^p a_t) \mu_k$. Let $\hat{r}_{i,k}$ denote the forecast of $r_{i,k}$, the h -step ahead forecast of the weights time series γ_k at instance i is then computed by

$$\hat{\gamma}_{i+h,k}^{AR} = s_{i+h,k} + d_{i+h,k} + \hat{r}_{i+h,k} \quad (14)$$

where the trend component d is computed by a linear extrapolation of the last five known trend values $d_{i-4,k}, \dots, d_{i,k}$ and the respective seasonal components s are obtained from the LOESS decomposition. The Akaike Information Criterion (AIC) is used to find the optimal model order p . We test the AIC for orders $p = 0, \dots, 6$ on all datasets. The AIC results were generally the best for $p = 1$ and we will continue to use the AR(1) model for all six datasets. A further interesting observation is that the dependency, measured by the pairwise correlation coefficients, among the remainder $r_{\cdot,k}$ time series after the decomposition slightly increases compared to the dependency between the original weights series $\gamma_{\cdot,k}$. Therefore, we will compare the univariate AR(1) model with the vector auto-regressive (VAR) model, see Luetkepohl (2005) and Box et al. (1994), on the joint remainder time series r . The VAR(1) model of the remainder $r(i) = (r_{i,1}, \dots, r_{i,3})$ has the following form

$$r(i) = \nu + A \cdot r(i-1) + u_i, \quad (15)$$

where $A \in \mathbb{R}^{K \times K}$ represents a fixed coefficient matrix, u a zero mean white noise process and ν the intercept. When μ denotes the mean of r , the intercept is equivalently defined as $\nu = (I - A)\mu$. Let $\hat{r} = (\hat{r}_{i,1}, \dots, \hat{r}_{i,K})$ denote the VAR forecast of r , the h -step ahead forecast of the weights time series γ at instance i is equivalent to the AR forecast computed by

$$\hat{\gamma}_{i+h,k}^{VAR} = s_{i+h,k} + d_{i+h,k} + \hat{r}_{i+h,k} \quad \text{for all } k = 1, \dots, K, \quad (16)$$

where the trend component d is again computed by a linear extrapolation of the last five known trend values. As in the comparison of both HW models, the forecast accuracy of the AR and the VAR forecast is tested on the sample of the 1 till 15 step ahead forecast, starting at instance 110 and computed on all datasets. The mean squared forecast errors ϵ_k in the

Dataset	ϵ_1 AR	ϵ_1 VAR	ϵ_2 AR	ϵ_2 VAR	ϵ_3 AR	ϵ_3 VAR
A-Ind	2.14	2.12	0.60	0.58	1.09	1.09
A-Acc	2.97	2.93	0.50	0.49	0.42	0.42
B-Ind	0.31	0.30	0.34	0.34	0.33	0.33
B-Acc	11.20	11.15	2.57	2.55	1.09	1.09
C-Ind	1.67	1.69	2.40	2.41	0.31	0.30
C-Acc	2.24	2.27	0.81	0.79	0.37	0.38

Table 3: AR and VAR mean squared errors between the actual and forecasted $\gamma_{.,k}$.

weights time series γ_k are shown in Table 3. Both forecasts generate roughly the same forecast errors, but a closer look shows that the VAR produces slightly smaller errors. More interesting is to compare the accuracy results of the Holt-Winters with the auto-regression forecasts, i.e., Table 2 with Table 3. The combination of seasonal-trend decomposition and auto-regression forecasts on the remainder increases the forecast accuracy significantly. In the remainder of the paper, we will work with the VAR(1) model as the second forecasting method as opposed to the AHW. The VAR forecast of the future booking horizon $\hat{x}_{i+h}^{VAR} = (\hat{x}_{i+h,1}^{VAR}, \dots, \hat{x}_{i+h,m}^{VAR})$ is obtained by

$$\hat{x}_{i+h}^{VAR} = \hat{\gamma}_{i+h,1}^{VAR} \cdot f_1 + \dots + \hat{\gamma}_{i+h,K}^{VAR} \cdot f_K. \quad (17)$$

4. Forecast Updating

With the previously described method, the forecaster is able to compute a forecast of an entire booking horizon, i.e., a forecast of the accumulated booking curve or the estimated incoming reservations for each day in the booking horizon. The forecasting methods work only on completed booking horizons. This means that we are not updating a forecast for future weeks when the realization of week 1,2 or 3 become known. Furthermore, the information of the booking realizations for 1,2 or 3 weeks prior to the forecast date are not used in the computation, because their booking horizon is not completed yet.

Therefore we propose the following general updating procedure, which includes one of the forthcoming updating methods. The procedure is as follows:

1. Forecast the next booking horizon based on data of all completed product instances (of which all realizations are known).
2. If realizations of the forecasted booking process are known, update the future part of the horizon accordingly.

3. If a further forecast is required, regard the forecasted (and updated) horizon as completed and go back to 1.

The application of the procedure to our datasets (being at arrival of product instance i) is as follows:

- (1) A one week ahead forecast of the complete booking horizon for instance $i + 1$ (the next week arrival) is generated based on data of all completed booking horizons (instances $1, \dots, i$). Realizations of the first three weeks in the booking horizon are already known and the forecast of reservations in the last week for instance $i + 1$ are adjusted based on this.
- (2) The adjusted forecast of $i + 1$ is further regarded as the completed booking horizon of instance $i + 1$.
- (3) A one week ahead forecast of the complete booking horizon for instance $i + 2$ is computed based on the data of all instances $1, \dots, i + 1$. Realizations of the first two weeks are already known and the forecast for the last two weeks of the booking horizon is adjusted.
- (4) The adjusted forecast of $i + 2$ is further regarded as the completed booking horizon of instance $i + 2$.
- (5) A one week ahead forecast of the complete booking horizon for instance $i + 3$ is computed based on the data of all instances $1, \dots, i + 2$. The realization of the first week is already known and the forecast for the following three weeks is adjusted.
- (6) The adjusted forecast of $i + 3$ is further regarded as the completed booking horizon of instance $i + 3$.
- (7) Finally, a one week ahead forecast of the complete booking horizon for instance $i + 4$ is computed based on the data of all instances $1, \dots, i + 3$. Because we are considering a booking horizon of four weeks, no bookings for this instance are known and the forecast can not be adjusted.

In the following we will discuss two forecast updating methods, as described in Shen and Huang (2008).

Let us concentrate on instance $i + 1$, the forecast of (11) or (17) can be written as

$$\hat{x}_{i+1} = F\hat{\gamma}_{i+1}, \quad (18)$$

where $F = (f_1, \dots, f_K)$ denotes a $m \times K$ matrix formed by the base vectors and $\hat{\gamma}_{i+1} =$

$(\hat{\gamma}_{i+1,1}, \dots, \hat{\gamma}_{i+1,K})^T$ represents a column vector. Let the superscript a denote that we only consider the first a columns of a matrix or components of a vector. When x_{i+1}^a becomes known, we can compute the forecast error ϵ_{i+1}^a by

$$\epsilon_{i+1}^a = x_{i+1}^a - \hat{x}_{i+1}^a = x_{i+1}^a - F^a \hat{\gamma}_{i+1}. \quad (19)$$

The direct least squares (LS) method would now try to solve the problem

$$\hat{\gamma}_{i+1}^{LS} = \operatorname{argmin}_{\hat{\gamma}_{i+1}} \|\epsilon_{i+1}^a\|^2, \quad (20)$$

to find the γ_{i+1} values for which the forecast of the first a days fits the actual bookings x_{i+1}^a best. The LS solution can be obtained by

$$\hat{\gamma}_{i+1}^{LS} = ((F^a)^T F^a)^{-1} (F^a)^T x_{i+1}^a. \quad (21)$$

To uniquely define $\hat{\gamma}_{i+1}^{LS}$, we need of course that $a \geq K$. In our case $K = 3$ or 5 , the booking horizon is further in days and the forecast updates are made weekly, $a = 7, 14$ and 21 . The idea is to apply the solution of (20) in (18) to obtain the direct least squares forecast update \hat{x}_{i+1}^{LS} by

$$\hat{x}_{i+1}^{LS} = F \hat{\gamma}_{i+1}^{LS}. \quad (22)$$

Clearly this is a very volatile updating method and the forecast update will not be too reliable for small a values compared to m , length of the whole booking horizon. Therefore we suggest the penalized least squares method (PLS), which works as the LS method but it penalizes large deviations from the original time series (TS) forecast. The optimization problem (20) is altered with the parameter λ to

$$\hat{\gamma}_{i+1}^{PLS} = \operatorname{argmin}_{\hat{\gamma}_{i+1}} \|\epsilon_{i+1}^a\|^2 + \lambda \|\hat{\gamma}_{i+1} - \hat{\gamma}_{i+1}^{TS}\|^2, \quad (23)$$

where $\hat{\gamma}^{TS}$ denotes the original time series forecast. We observe that if $\lambda = 0$, $\hat{\gamma}^{PLS} = \hat{\gamma}^{LS}$, and for $\lambda \rightarrow \infty$, $\hat{\gamma}^{PLS} = \hat{\gamma}^{TS}$. As shown in Shen and Huang (2008), the PLS updated forecast can be computed with

$$\hat{\gamma}_{i+1}^{PLS} = ((F^a)^T F^a + \lambda I)^{-1} ((F^a)^T x_{i+1}^a + \lambda \hat{\gamma}_{i+1}^{TS}). \quad (24)$$

And finally, the PLS updated forecast of the future booking horizon \hat{x}_{i+1}^{PLS} is obtained by

$$\hat{x}_{i+1}^{PLS} = \hat{\gamma}_{i+1,1}^{PLS} f_1 + \dots + \hat{\gamma}_{i+1,K}^{PLS} f_K. \quad (25)$$

One other more intuitive updating approach is the historical proportion (HP) method. The accuracy of the forecast is simply computed by the ratio of already observed realization and their forecasted values. Suppose we are at updating point a , i.e., realizations of the first a days in the booking horizon are known. The ratio R is given by

$$R = \frac{\sum_{j=1}^a x_{i+1,j}}{\sum_{j=1}^a \hat{x}_{i+1,j}}, \quad (26)$$

keeping in mind that \hat{x}_{i+1} denotes the time series based forecast of x_{i+1} . The HP updated forecast for the remaining booking days is the with R scaled \hat{x}_{i+1} ,

$$\hat{x}_{i+1,j}^{\text{HP}} = R \cdot \hat{x}_{i+1,j} \quad j = a + 1, \dots, m. \quad (27)$$

In the following section we will compare the PLS and HP updating method with the forecast results that are not updated.

5. Numerical Results

In this section we will compare all combinations of the additive Holt-Winters and the vector auto-regressive forecasts with the two previously proposed updating methods penalized least squares and historical proportion, as well as with the not updated forecasts (NU). The number behind the abbreviation of the forecasting method (AHW or VAR) denotes the number of base vectors K used in the singular value decomposition. In our test case we are working with $K = 3$ or 5 . The evaluation set consists of the last 22 instances, i.e., arrival weeks, of our six datasets (instances 134-155). Thus, the evaluation is made in a time frame of five months. As measures of forecast accuracy the mean squared error (MSE) and the mean week relative absolute error (MWRAE) are computed for the four booking horizon weeks. The squared error (SE) and the week relative absolute error (WRAE) are defined for instance i and weeks $w = 1, \dots, 4$ by

$$\begin{aligned} \text{SE}(i, w = 1) &= \sum_{k=1}^7 (x_{i,k} - \hat{x}_{i,k})^2 & \text{WRAE}(i, w = 1) &= \frac{|\sum_{k=1}^7 x_{i,k} - \sum_{k=1}^7 \hat{x}_{i,k}|}{\sum_{k=1}^7 x_{i,k}} \\ \text{SE}(i, w = 2) &= \sum_{k=8}^{14} (x_{i,k} - \hat{x}_{i,k})^2 & \text{WRAE}(i, w = 2) &= \frac{|\sum_{k=8}^{14} x_{i,k} - \sum_{k=8}^{14} \hat{x}_{i,k}|}{\sum_{k=8}^{14} x_{i,k}} \\ \text{SE}(i, w = 3) &= \sum_{k=15}^{21} (x_{i,k} - \hat{x}_{i,k})^2 & \text{WRAE}(i, w = 3) &= \frac{|\sum_{k=15}^{21} x_{i,k} - \sum_{k=15}^{21} \hat{x}_{i,k}|}{\sum_{k=15}^{21} x_{i,k}} \\ \text{SE}(i, w = 4) &= \sum_{k=22}^{28} (x_{i,k} - \hat{x}_{i,k})^2 & \text{WRAE}(i, w = 4) &= \frac{|\sum_{k=22}^{28} x_{i,k} - \sum_{k=22}^{28} \hat{x}_{i,k}|}{\sum_{k=22}^{28} x_{i,k}}. \end{aligned}$$

The MSE and the MWRAE are computed by averaging the SE and WRAE over our 22 evaluation instances. Both error measures are the ones used by the company, Bookit. The MSE gives insight into the accuracy on daily level, while the MWRAE provides the proportional absolute difference in week totals. The optimal λ parameter for the PLS are found by mini-

Forecast	Dataset	λ_1	λ_2	λ_3	Dataset	λ_1	λ_2	λ_3
VAR 3	A-Acc	0.0114	0.0452	0.1691	A-Ind	0.1676	0.2893	2.0623
VAR 5		0.0040	0.0182	0.0615		0.2617	0.8529	0.9294
AHW 3		0.0029	0.0213	0.0411		0.0736	0.1117	0.4801
AHW 5		0.0010	0.0072	0		0.1511	0.6909	0.3270
VAR 3	B-Acc	0.0206	0.0089	0.0594	B-Ind	148.6172	0	0.2699
VAR 5		0.0090	0.0093	0.0289		177.9697	0.0920	0.3775
AHW 3		0.0141	0.0099	0.0322		0.5054	0.1471	0.2225
AHW 5		1.1320	2.0711	0.0123		0.4020	0.3317	0.2757
VAR 3	C-Acc	0.0276	0.0452	0.0962	C-Ind	0.0236	0.4586	0.0635
VAR 5		0.0084	0.0275	0.0583		11.1344	0.2037	0.1561
AHW 3		0.0033	0	0.0015		0.0421	0	0.1801
AHW 5		0.0011	0	0.0187		0.0761	0	0.1431

Table 4: λ parameter for PLS updating methods, respectively for dataset and booking week.

mizing the MSE updating error at the last 22 instances of the testing and estimation sample (instances 109 till 130), see Table 4 for the final values. λ_1, λ_2 and λ_3 are respectively used in the updating in booking horizon weeks 1,2 and 3. The best forecast and updating method combinations for each dataset, which minimize the forecast error per booking horizon week, are given in Table 5. All generated MSE and MWRAE error results are shown in Table 6, the smallest errors for dataset and booking week combination are highlighted by an asterisk (*).

For the MSE we find in weeks 2, 3 and 4 the smallest values for the PLS updated forecasts. Except for the C-Ind dataset in week 4, there the VAR 5 PLS value exceeds the VAR 3 HP value by 12, which only corresponds to a mean absolute daily error of 0.5. We also observe that the VAR outperforms the AHW forecast, except for the B-Acc dataset in week 1 and 2, but here we find again an insignificant increase of the MSE by only 3 and 7 compared with the VAR 5 PLS results. Considering the MWRAE we initially observe that the accuracy increases with the amount of reservations contained in a dataset. Consequently the C datasets have the lowest MWRAE, as they also hold the most reservations. The VAR 5 forecast of week 4 in C-Acc without updating already has a MWRAE of 0.18, but can still be decreased to 0.08 by the PLS updating method. Looking at the B-Acc dataset, again week 4 and VAR 5, the

Dataset	MSE			
	1 st week	2 nd week	3 rd week	4 th week
A-Ind	VAR 5 NU & PLS	VAR 3 PLS	VAR 3 PLS	VAR 3 PLS
A-Acc	VAR 3 NU	VAR 3 PLS	VAR 3 PLS	VAR 3 PLS
B-Ind	VAR 3 NU & PLS	VAR 3 NU & PLS	VAR 3 PLS	VAR 3 PLS
B-Acc	AHW 3 NU	AHW 3 PLS	VAR 3 PLS	VAR 3 PLS
C-Ind	VAR 3 HP & PLS	VAR 3 NU & 5 PLS	VAR 3&5 PLS	VAR 3 HP
C-Acc	VAR 3 NU	VAR 3 PLS	VAR 5 PLS	VAR 5 PLS
Dataset	MWRAE			
	1 st week	2 nd week	3 rd week	4 th week
A-Ind	VAR 3 NU	VAR 3 PLS	AHW 3 PLS	AHW 5 PLS
A-Acc	VAR 3 NU	VAR 3 PLS	VAR 5 PLS	VAR 3 PLS
B-Ind	VAR 3 PLS	VAR 3 PLS	VAR 3 HP	VAR 3 NU
B-Acc	AHW 3&5 NU	AHW 3 HP	VAR 3 PLS	VAR 5 PLS
C-Ind	VAR 3 HP & PLS	VAR 3 NU	VAR 5 NU	AHW 3 HP
C-Acc	VAR 5 NU & PLS	VAR 3 PLS	VAR 5 PLS	VAR 5 PLS

Table 5: Best forecast and updating methods respectively to minimization of MSE or MWRAE per booking week and dataset.

not updated forecast has a MWRAE of 0.73 which can be significantly decreased by the PLS updating to 0.2. At first glance at Table 5 in the MWRAE area, the PLS updating is still the best, but with less dominance compared to the MSE part. With a closer look at Table 6, we find that the PLS error values are very close to the best performing methods; compare for example the VAR 3 PLS values for A-Ind, B-Ind and C-Ind. Note that approximately 80% of all bookings are made within this last three weeks of the booking horizon and still 60% within the last two booking weeks. Therefore a forecast accuracy increase in the later part of the booking horizon is more important than in the early stages. Comparing the two forecasting methods additive Holt-Winters and vector auto-regression, we observe that the mean values of the VAR outperform the AHW forecast. This shows that the correlation between the base vectors should not be neglected. The results of the different updating methods are graphically illustrated in Figure 8, for one instance of the evaluation set for datasets A-Ind and A-Acc and base forecast VAR 3. For practitioners, the accurate forecast of the overall number of reservations is as important as the forecast of the booking curve or individual reservations per booking day. The total relative absolute error (TRAЕ) for “Ind” datasets is then defined by

$$\text{TRAЕ}(i) = \frac{|\sum_{k=1}^{28} \hat{x}_{i,k} - \sum_{k=1}^{28} x_{i,k}|}{\sum_{k=1}^{28} x_{i,k}}.$$

			MSE				MWRAE			
			1 st week	2 nd week	3 rd week	4 th week	1 st week	2 nd week	3 rd week	4 th week
A-Ind	VAR 3	NU	16	23	35	64	0.26*	0.32	0.45	0.37
		PLS	16	24*	31*	58*	0.28	0.27*	0.38	0.34
		HP	17	29	32	74	0.36	0.40	0.37	0.46
	VAR 5	NU	15*	25	38	66	0.30	0.32	0.46	0.37
		PLS	15*	27	34	58	0.34	0.29	0.39	0.35
		HP	16	30	38	75	0.34	0.42	0.40	0.46
	AHW 3	NU	32	47	47	98	1.07	0.91	0.65	0.45
		PLS	26	27	29	67	0.78	0.29	0.34*	0.34
		HP	26	58	30	205	0.78	0.52	0.35	0.66
	AHW 5	NU	34	50	51	101	1.05	0.97	0.65	0.44
		PLS	36	31	44	64	0.90	0.35	0.49	0.33*
		HP	36	54	37	295	0.90	0.55	0.36	0.68
A-Acc	VAR 3	NU	85*	282	929	2,047	0.45*	0.32	0.32	0.25
		PLS	89	137*	267*	334*	0.47	0.15*	0.14	0.07*
		HP	193	1,148	1,467	2,696	0.60	0.54	0.30	0.24
	VAR 5	NU	100	284	873	2,206	0.50	0.31	0.31	0.26
		PLS	135	288	272	404	0.56	0.19	0.13*	0.08
		HP	167	1,094	1,443	2,643	0.59	0.54	0.30	0.22
	AHW 3	NU	373	1861	6,189	19,527	1.27	0.93	0.84	0.67
		PLS	325	642	341	517	0.99	0.27	0.15	0.09
		HP	325	16,320	5,203	207,650	0.99	0.74	0.41	0.67
	AHW 5	NU	353	2,029	6,454	18,760	1.21	0.98	0.86	0.64
		PLS	345	561	477	785	0.98	0.27	0.18	0.11
		HP	345	34,193	5,153	100,050	0.98	0.89	0.41	0.57
B-Ind	VAR 3	NU	1.5*	3.1*	5.4	2.7	0.61	0.69	0.74	0.50*
		PLS	1.5*	3.1*	5.3*	2.6*	0.60*	0.67*	0.61	0.54
		HP	1.5*	3.4	5.4	2.7	0.62	0.95	0.56*	0.59
	VAR 5	NU	1.6	3.1	5.4	2.9	0.62	0.73	0.66	0.62
		PLS	1.6	3.2	5.4	2.7	0.63	0.75	0.77	0.62
		HP	1.6	3.4	5.6	2.9	0.65	0.94	0.73	0.70
	AHW 3	NU	1.6	3.6	5.5	3.1	0.72	0.88	0.77	0.69
		PLS	1.5*	3.5	5.3*	2.7	0.75	0.69	0.68	0.63
		HP	1.5*	3.7	5.4	2.9	0.75	0.93	0.72	0.71
	AHW 5	NU	1.7	3.6	5.5	3.3	0.76	0.94	0.73	0.77
		PLS	1.5*	3.7	5.3*	2.9	0.77	0.74	0.86	0.73
		HP	1.5*	3.7	5.4	4.6	0.77	0.94	0.74	1.07
B-Acc	VAR 3	NU	13	47	85	93	3.37	3.71	1.65	0.74
		PLS	13	41	43*	33*	3.01	1.89	0.76*	0.28
		HP	14	282	764	434	3.47	0.86	1.11	1.11
	VAR 5	NU	13	54	75	92	3.49	3.85	1.54	0.73
		PLS	12	36	75	43	2.99	2.53	1.10	0.20*
		HP	15	284	1,638	361	3.90	0.87	1.23	0.80
	AHW 3	NU	9*	36	87	128	1.97*	3.70	3.23	1.08
		PLS	14	29*	58	48	3.65	1.50	1.95	0.32
		HP	14	39	574	1245	3.65	0.68*	1.15	0.51
	AHW 5	NU	10	36	92	125	1.97*	3.63	3.43	1.06
		PLS	15	36	201	63	3.68	3.43	6.60	0.27
		HP	15	39	492	205	3.68	0.71	1.11	0.55
C-Ind	VAR 3	NU	33	25*	43	85	0.42	0.20*	0.25	0.37
		PLS	32*	27	37*	96	0.38*	0.29	0.23	0.38
		HP	32*	31	46	80*	0.38*	0.34	0.27	0.37
	VAR 5	NU	37	26	43	87	0.51	0.23	0.21*	0.38
		PLS	35	25*	37*	92	0.47	0.21	0.25	0.40
		HP	35	35	49	81	0.45	0.38	0.30	0.38
	AHW 3	NU	52	65	122	265	0.81	0.62	0.72	0.76
		PLS	41	26	43	109	0.70	0.26	0.28	0.47
		HP	41	33	42	94	0.70	0.35	0.27	0.34*
	AHW 5	NU	62	67	124	281	0.93	0.64	0.74	0.77
		PLS	49	42	61	136	0.77	0.39	0.39	0.59
		HP	49	48	40	99	0.77	0.47	0.23	0.36
C-Acc	VAR 3	NU	156*	437	886	2,629	0.50	0.25	0.18	0.17
		PLS	160	112*	371	963	0.47	0.09*	0.09	0.09
		HP	233	2,513	2,468	4,900	0.54	0.44	0.29	0.25
	VAR 5	NU	180	505	990	2,920	0.46*	0.27	0.19	0.18
		PLS	188	128	286*	723*	0.46*	0.11	0.08*	0.08*
		HP	255	2,781	2,762	5,981	0.56	0.43	0.29	0.25
	AHW 3	NU	728	4,236	12,444	31,035	0.96	0.82	0.78	0.75
		PLS	445	646	573	996	0.75	0.20	0.10	0.10
		HP	445	4,924	1,286	9,693	0.75	0.57	0.22	0.30
	AHW 5	NU	831	4,840	12,116	27,654	0.98	0.90	0.77	0.70
		PLS	505	564	481	770	0.76	0.26	0.11	0.08
		HP	505	5,622	1,149	15,958	0.76	0.63	0.21	0.30

Table 6: Mean squared errors (MSE) and mean week relative absolute errors (MWRAE) for all datasets and forecast combinations - the asterisk highlights the smallest errors for each dataset and booking week combination.

In case of accumulated “Acc” datasets the TRAE is simply computed by the last column values

$$\text{TRAE}(i) = \frac{|\hat{x}_{i,28} - x_{i,28}|}{x_{i,28}}$$




Figure 8: Example of reservation actuals vs. updated and not updated forecasts for one instance of datasets A-Acc and A-Ind.

In Table 7 the different mean and median TRAE values are shown for all datasets and updating methods applied to a VAR 3 forecast. The PLS is completely outperforming the

Dataset	Mean TRAE			Median TRAE		
	NU	PLS	HP	NU	PLS	HP
A-Ind	0.28	0.21	0.27	0.24	0.16	0.28
B-Ind	0.46	0.32	0.40	0.44	0.33	0.46
C-Ind	0.22	0.17	0.21	0.23	0.16	0.20
A-Acc	0.22	0.11	0.28	0.21	0.09	0.24
B-Acc	0.54	0.30	0.92	0.41	0.26	0.66
C-Acc	0.18	0.11	0.27	0.16	0.10	0.25

Table 7: Mean and median of total relative absolute error (TRAE) for all updating methods and datasets with VAR 3 as base forecast.

not updating and HP updating in all cases. Further, the TRAE for datasets of the same region is always minimized on the accumulated dataset. Two interesting observations are: First, the HP updating is only better than not updating in the case of “Ind” datasets, and second that

PLS updating on “Ind” datasets results in lower TRAE than not updating on “Acc” datasets. Finally, we are interested in testing the sensitivity of the PLS to its λ values. Therefore, we

	MSE			
	1 st week	2 nd week	3 rd week	4 th week
A-Ind λ_{opt}	16	22	29	58
A-Ind λ_{est}	16	24	31	58
A-Acc λ_{opt}	98	99	219	338
A-Acc λ_{est}	89	137	267	334
	MWRAE			
	1 st week	2 nd week	3 rd week	4 th week
A-Ind λ_{opt}	0.27	0.25	0.33	0.35
A-Ind λ_{est}	0.28	0.27	0.38	0.34
A-Acc λ_{opt}	0.47	0.14	0.11	0.07
A-Acc λ_{est}	0.47	0.15	0.14	0.07

Table 8: MSE and MWRAE values of VAR 3 PLS with optimal (opt) and estimated (est) λ values.

compare the MSE and MWRAE generated using λ values estimated on historical data versus the error values generated by optimal λ values, i.e., estimated by minimizing the MSE over the evaluation set. In the following the VAR 3 is used as base forecast method. The optimal λ_1, λ_2 and λ_3 values are 0.4991, 0.0841 and 5.5078 for A-Ind and 0.0372, 0.0039 and 0.1789 for A-Acc, which are deviating from the so far used values (see Table 4). The generated MSE and MWRAE results of the optimal λ 's are given in Table 8. As expected, using the optimal λ in the PLS updating results in an accuracy increase. But, we also find the error decrease by using the optimal lambda values to be considerably small. These results support our positive validation on the robustness of the PLS method.

6. Conclusion

In this study we have analyzed forecast updating methods using actual hotel reservation data from three different regions. Statistical tests have shown that there is a significant correlation between early and late bookings. Consequently not updating the demand forecast when early realizations in the booking horizon become available means ignoring important information that can dramatically affect the end result. But, also for datasets with low correlation between early and late bookings (region B in our data set) we observe a significant accuracy increase by updating.

The forecast updating is performed dynamically when new demand realizations become

available, in our case weekly. The initial forecast results are then updated using either the penalized least squares (PLS) or the historical proportion method. In the case of a multi step ahead forecast, the base forecast produces multiple one step ahead forecasts on historical data and previous updated forecasts. We find that dynamic updating reservation forecasts using PLS is very beneficial in most situations with low and high correlation between different parts of the booking horizon, and is never significantly harmful compared to not updating. Also computationally the method is very fast and therefore feasible for use by practitioners in larger forecasting problems.

Singular value decomposition is applied to reduce the dimensionality of the forecasting problem and the results show its effectiveness. As base forecasts we use a multivariate vector autoregressive model and a univariate Holt-Winters model on the reduced forecasting problem. The results show that the VAR outperform the AHW. Thus, the dependency between base vectors after the SVD should not be ignored. In addition, an increase of base vectors from 3 to 5 generally does not result in lower error values. Overall, the VAR 3 forecast method seems to be the best base forecast for our datasets.

So far in this study we have assumed the demand to be independent for different hotel products, i.e., combinations of region, hotel category, weekday of arrival and length of stay. In reality a buying decision depends crucially on the offered alternatives. Therefore the application of updating methods to choice based demand forecasts should also be investigated. As a result of the case study, Bookit B.V. decided to embed a dynamic updating method in their forecasting system.

7. Acknowledgements

The authors thank Bookit B.V. for the opportunity to perform the case study on real hotel reservation data. Further thanks goes to the handling editor Rob Hyndman and two anonymous referees for their valuable comments that considerably improved the paper.

References

- Box, G.E.P., Jenkins, G.M., and Reinsel, G.C. (1994). *Time Series Analysis: Forecasting and Control*, New Jersey: Prentice Hall, Third Edition.
- Chiang, W.C., Chen, J.C.H., and Xu, X. (2007). An Overview of Research on Revenue Management: Current Issues and Future Research. *International Journal of Revenue Management*, Vol. 1, No. 1, pp. 97-128.
- Cleveland, R.B., Cleveland, W.S., McRae, J.E., and Terpenning, I. (1990). STL: A Seasonal-Trend Decomposition Procedure Based on Loess. *Journal of Official Statistics*, Vol. 6, No. 1, pp. 3-73.
- O'Connor, M., Remus, W., and Griggs, K. (2000). Does Updating Judgmental Forecasts Improve Forecast Accuracy? *International Journal of Forecasting*, Vol. 16, pp. 101-109.
- Haensel, A. and Koole, G.M. (2010). Estimating Unconstrained Demand Rate Functions using Customer Choice Sets. *Journal of Revenue and Pricing Management* advance online publication, April 30, 2010 (doi:10.1057/rpm.2010.1).
- Holt, C.C. (1957). Forecasting Seasonals and Trends by Exponentially Weighted Moving Averages. *ONR Research Memorandum Carnegie Institute* 52.
- Liu, P.H., Smith, S., Orkin, E.B., and Carey, G. (2002). Estimating Unconstrained Hotel Demand Based on Censored Booking Data. *Journal of Revenue and Pricing Management*, Vol. 1, No. 2, pp. 121-138.
- Lütkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis*, Berlin: Springer.
- Pölt, S. (1998). Forecasting is difficult especially if it refers to the future. *AGIFORS - Reservations and Yield Management Study Group Meeting Proceedings*.
- Rajopadhye, M., Ben Ghalia, M., Wang, P.P., Baker, T., and Eister, C.V. (2001). Forecasting Uncertain Hotel Room Demand. *Information Science*, Vol. 132, pp. 1-11.
- Shen, H., and Huang, J.Z. (2008). Interday Forecasting and Intraday Updating of Call Center Arrivals. *Manufacturing and Service Operations Management*, Vol. 10, No. 3, pp. 391-410.
- Talluri, K.T., and Van Ryzin, G.R. (2004a). Revenue Management Under a General Discrete Choice Model of Customer Behavior. *Management Science*, Vol. 50, No. 1, pp. 15-33.
- Talluri, K.T., and Van Ryzin, G.R. (2004b). *The Theory and Practice of Revenue Management*, New York: Springer.
- Van Ryzin, G.R. (2005). Future of Revenue Management - Models of Demand. *Journal of Revenue and Pricing Management*, Vol. 4, No. 2, pp. 204-210.
- Vinod, B. (2004). Unlocking the Value of Revenue Management in the Hotel Industry. *Journal of Revenue and Pricing Management*, Vol. 3, No. 2, pp. 178-190.
- Weinberg, J., Brown, L.D., and Stroud, J.R. (2007). Bayesian Forecasting of an Inhomogeneous Poisson Process with Applications to Call Center Data. *Journal of the American Statistical Association*, Vol. 102, No. 480, pp. 1185-1198, Applications and Case Studies.
- Weatherford, L.R., and Kimes, S.E. (2003). A Comparison of Forecasting Methods for Hotel Revenue Management. *International Journal of Forecasting*, Vol. 19, pp. 401-415.
- Winters, P.R. (1960). Forecasting Sales by Exponentially Weighted Moving Averages. *Management Science*, Vol. 6, pp. 324-342.
- Zakhary, A., Atiya, A.F., El-Shishiny, H., and Gayar, N.E. (2009). Forecasting Hotel Arrivals and Occupancy using Monte Carlo Simulation. *Journal of Revenue and Pricing Management* advance online publication, November 13, 2009 (doi:10.1057/rpm.2009.42).
- Zeni, R.H. (2007). Can we really look to the past to forecast future demand? *Journal of Revenue and Pricing Management*, Vol. 6, No. 4, pp. 312-314.