Mathematics and Computer Science
2021; 6(1): 8-15
http://www.sciencepublishinggroup.com/j/mcs
doi: 10.11648/j.mcs.20210601.12
ISSN: 2575-6036 (Print); ISSN: 2575-6028 (Online)

Patterns in Hotel Revenue Management Forecasting Systems: Improved Sample Sizes, Frozen Intervals, Horizon Lengths, and Accuracy Measures

Victor Pimentel[1], Aysajan Eziz[2], Tim Baker[3, *]

[1]Department of Management, College of Business, New Mexico State University, Las Cruces, the United States

[2]Ivey Business School, University of Western Ontario, London, Canada

[3]Carson College of Business, Washington State University, Richland, United States

Email address:
vip@nmsu.edu (V. Pimentel), aeziz@ivey.ca (A. Eziz), bakert@wsu.edu (T. Baker)

[*]Corresponding author

To cite this article:
Victor Pimentel, Aysajan Eziz, Tim Baker. Patterns in Hotel Revenue Management Forecasting Systems: Improved Sample Sizes, Frozen Intervals, Horizon Lengths, and Accuracy Measures. *Mathematics and Computer Science*. Vol. 6, No. 1, 2021, pp. 8-15.
doi: 10.11648/j.mcs.20210601.12

Received: January 6, 2021; **Accepted**: January 20, 2021; **Published**: January 28, 2021

Abstract: Research in hotel revenue management system design has not paid much attention to the demand forecasting side of the system. And the research that has examined forecasting has tended to focus on the comparison of specific forecaster methodologies, as opposed to prioritizing how a total system should be parameterized: how far in the future should projections be, how much data to use to update each specific parameter, which measure of forecast error to use, and how long to freeze each parameter/forecast before updating. This paper fills this prioritization void by utilizing a full-functionality hotel reservation system simulation validated by the revenue management staff of a major hotel chain as the basis for running screening experiments on an exhaustive set of forecaster parameters with regards to their impact on bottom-line system performance (average nightly net revenue, where net revenue refers to total room rate receipts minus an overbooking per person penalty that estimates the discounted lost sales of future revenues). A screening experiment is run for each general type of operating environment (demand intensity, degree of market segment differentiation) that a property might face. We find that only two parameters are significant: the final combined forecast horizon length and how long that final forecast is frozen before updating. We find that these two factors interact in a negative fashion to influence net revenue.

Keywords: Forecasting, Simulation, Capacity Analysis, Statistics

1. Introduction

According to Kummumkal ad Talluri, revenue management is the forecasting and allocation of perishable assets across assorted market segments to maximize expected revenue over a short-term planning horizon. [1] A perishable asset—a hotel room, rental car, airline seat, broadcast advertising time slot, theatre seat, cruise line space, for instance—is an asset whose value drops to zero instantly at some point in time. Rather than depreciating in value gradually, like many physical goods. For example, once a stayover night has passed, an unsold hotel room's lost revenue can never be recaptured. Revenue maximization is the objective since in the short-run the associated cost structures (e.g., hotel physical configuration, salaried and hourly employees to serve guests) cannot be changed.

The bottom-line benefits of revenue management are substantial. In Zhang and Weatherford's work, a new optimization algorithm that involved forecasting would generate an estimated $10 million in additional revenue. [2] Druckman reported 8.7% increase in revenues due to installing a revenue management system for rental properties has been made at one property management system. [3]

Assorted algorithms have been proposed to implement the goals of revenue management. Vives *et al.* in the resort industry emphasized the importance of demand forecasting.

[4] Rather, according to Duffy forecasts of market segment demand for these assets are done over the planning horizon as inputs to the subsequent allocation algorithms. [5] A stochastic network optimization was developed by Lai and Ng and is a common approach to allocate the various hotel rooms to customers with different price sensitivities (i.e., rate classes), arrival dates, and lengths-of-stay. [6] Where the objective is to maximize expected revenue (net of cancellations and no-shows) over some planning horizon. Such an optimization can be modified to not only do the allocation across the assorted rate classes, but to also to approximately overbook the hotel to account for cancellations and no-shows.

The approaches to forecasting demand has received much less attention in the literature. And much of this forecasting literature has compared-contrasted assorted quantitative forecasting techniques, rather than how a given set of techniques should best be implemented. For instance, Weatherford and Kimes evaluated assorted demand forecasting methods on Marriott data—exponential smoothing, moving average, regression approaches, booking profile methods and combination forecasts. [7] But they did not closely examine how their best methods (exponential smoothing, moving average) should be implemented in terms of (a) how much data to use in parameter estimation, (b) how long to freeze the forecasts before updating them, (c) how long should the forecasting horizon be, and (d) what metrics to use to evaluate forecast accuracy (e.g., root mean square error to penalize large errors versus mean absolute deviation).

Our paper focuses on the best forecasting implementation for (a)-(c) using the current industry standard forecasting techniques from Duffy. [5] We do this by developing a complete simulation model of a hotel revenue management system: customer reservation patterns, uncensoring of observed demand, demand forecasting of uncensored demand, overbooking and optimization simultaneously, acceptance/rejection of individual reservation requests based on the overbooking-optimization outputs. All of the components of this system reflect best industry practice today, except that our simultaneous overbooking-allocation algorithm extends best practice in Duffy by doing these two activities jointly. [5]

We use a Plackett-Burman screening approach to determine the most impactful forecast implementation factors on the revenue-generating capability of this system.

The only two articles that deal with the optimal implementation of revenue management forecasting methods are Weatherford and Belobaba and Weatherford *et al.* [8, 9]. Weatherford and Belobaba were focused on airline data [8]. They compared the revenue impact of the accuracy of the forecast point estimate relative to the accuracy in estimating forecast error variability. But they did not evaluate the amount of data for estimating parameters, length of frozen period, or measure-of-merit used when updating parameters, or forecast horizon length. Likewise Weatherford *et al.* also did not consider these four factors that we are evaluating in this paper. [9]

2. Materials and Methods

Azadeh *et al.* have developed a taxonomy of assorted methods used in practice to estimate "true demand", as opposed to "observed demand"). [10] Our method used matches that used in practice today per Duffy, so it can be thought of as a control variable. [5] The core of our unconstraining method is implemented in ~~a~~ the choice model of van Ryzin and Vulcano.[11]. Yukel has augmented the forecasting techniques that we use with a Delphi method override of forecasts in a real time hotel environment in Ankara. [12] For the purposes of seeing patterns in optimal data amount, frozen period, horizon length, and measure-of-merit, omitting user overrides acts as another control. Weatherford and Belobaba determined that forecast point estimates are more important than limiting forecast error sizes. [8] Weatherford *et al.* determined that in a hotel environment disaggregated forecasts by rate class and duration are more accurate than aggregating over these dimensions. [9] Our paper always works with the disaggregated data, except in the overbooking algorithm.

Our overbooking method is contained in Huang *et al.* that is in an airline multi-leg context. [13] We aggregate all rate class-duration-arrival date combinations that cover a given control date. Phumchurri and Maneesephon have developed a closed-form dynamic programming solution to at most a two rate class hotel overbooking problem. [14] But real world overbooking situations often involve eight or more rate classes according to Duffy. [5]

Our allocation algorithm has the same functionality as Lai and Ng. [6] Except that we use a nonlinear programming solver rather than a network solver. Fouad *et al.* have developed an optimization in simulation approach to overbooking like we have, except that they ignore the allocation portion. [15]

Koch does integrated overbooking-allocation, but their context is choice modeling. [16]

We develop a simulation of a hotel reservation system, with revenue management system embedded, that accounts for all of the reservation and revenue management activities in practice according to Duffy, except for group bookings. [5] We have simply scaled it down to two rate classes, a three day maximum length-of-stay, and a fourteen day maximum booking horizon to permit extensive computational experiments for this study.

One replication of the simulation will be 113 nights, with the first 21 nights of revenue ignored since initial condition bias was present then.

The overall essence of this simulation is, on each day, to (1) perform true net demand forecasts for the planning horizon from historical estimated true demand, (2) run genetic algorithm to simultaneously do overbooking and allocation—the ultimate output are the bid prices, (3) generate true demand (see Pimentel *et al.* for terminology) from the Poisson demand processes, (4) use the bid prices to accept/reject each reservation request, (5) convert the reservations that get accepted to estimated historical true demand by considering how often each rate class-duration-arrival date combination was open for sale, (6) convert the estimated historical true demand to estimated historical true net demand by applying cancellation and no-show probabilities, (7) count revenue and overbooking penalties for the night, (8) return to step (1) unless the 113 night simulation window is hit. [17]

Here are the complete functional details of the simulation, listed step-buy-step:

Step 1: Initialize the system. First, the hotel has to be booked. The sum of the Poisson true demands (not estimated true demand) at the bottom of The data tables in Pimentel *et al.* should be used, filling up the hotel by highest rate class first, followed by longest to shortest duration. [18]. The Holt-Winter's trend term should be set to 0 for all rate class-duration-arrival date combinations. The Holt-Winter's base term should be set to the mean weekly true demand from Pimentel *et al.*, minus cancellation and no-show probabilities, for each rate class-duration-arrival date combination [18]. The Holt-Winter's seasonal indices should all be set to 1 for each rate class-duration-arrival date combination. The initial data value in the booking profile forecast should be the top intensity from [18]. The booking profiles for each rate class-duration-arrival date combination should be based on the Pimentel *et al.* cumulative Poisson intensities. The weights assigned to the booking profile forecast and Holt-Winter's forecast to make up the final combined forecast should be 0.5 and 0.5. The shape and scale parameters of the gamma forecast error probability density function should be derived via the method of moments assuming a coefficient of variation of $\frac{1}{\sqrt{\mu}}$, where μ is the sum of the Poisson intensities across all days until first stayover from Pimentel *et al.* for a given rate class-duration-arrival date, and μ is the mean. The first day of the simulation is Monday. [17]

Step 2: Based on the most recent true net demand estimates, perform the Holt-Winter's true net demand forecast for each given arrival day-of-the-week out until the end of the planning horizon for each rate class-duration combination. Note that estimated true demand equals actual true demand at simulation initialization.

Step 3: Based on the most recent true net demand, perform the booking profile true net demand forecast for arrivals for the next day, the following day,…, the final day in the planning horizon. For each rate class-duration combination.

Step 4: For each booking profile true net demand forecast in step 3, execute the corresponding final true net demand forecast by doing a linear combination of the Holt-Winter's forecast and booking profile forecast.

Step 5: Execute the genetic algorithm from Pimentel *et al.* Given the most recent true net demand forecasts across the planning horizon. The outputs of this algorithm are the bid prices. [17]

Step 6: For each cell entry in Pimentel *et al.*, sample from the Poisson distribution using the intensity as the mean to obtain true demands for each rate class-duration-arrival date combination across the planning horizon. [17]

Step 7: For each true demand from Step 6, accept it as a booking if and only if the bid price is less than or equal to the sum of the corresponding nightly revenues from Baker and Collier. [18]

Step 8: For the historical bookings time series by arrival day-of-week for each rate class-duration combination, divide by the fraction of the time across all reading periods that this rate class-duration-arrival date combination was open for sale (i.e, the sum of bid prices was less than or equal to the sum of nightly revenues). This will give you the estimated true demand.

Step 9: For these estimated true demands, now multiply by (1 – cancellation rate)**(planning horizon length)*(1 – no-show rate) to obtain estimated true net demand.

Step 10: Given the accepted bookings that would like to stayover on this night, book as many as possible who arrive on this night until physical capacity is reached. Count the total night's revenue, and compute net revenue by subtracting the number overbooked times the overbooking penalty.

Step 11: Re-optimize the parameters of the Holt-Winter's forecaster and their relative weights in the final forecast. Do this reoptimization based on minimizing either RMSE or MAD from prior forecasts via enumeration.

Step 12: Recalculate the shape-scale parameters of the forecast error distribution. Recalculate the booking profile.

Step 13: If we are at day 113, stop—this simulation replication has been completed. Otherwise, move to the next day and go to Step 2.

3. Calculation

We wish to do an exploratory study to determine (1) patterns in the optimal number of historical data points to use in updating assorted forecaster parameters, (2) patterns in the optimal amount of time to freeze (i.e., leave unchanged) various forecaster parameters, (3) patterns in the best measure of merit for updating assorted forecaster parameters, and (4) patterns in how long a forecast horizon should be. By "optimal", we mean finding the best number of historical data points-frozen intervals-horizon lengths to maximize the expected net revenue. Thus, forecaster accuracy, measured by root mean squared error or mean absolute deviation, is not the ultimate goal in optimizing the design of the forecasting system; rather, the ultimate measure of merit is the average net revenue that the design will generate across a number of independent simulation replications. This is the most relevant measure of merit for designing a forecasting system since it is the "bottom-line" for any revenue management system. However, within the forecast parameter updating system, a more focused measure must be used for updating the parameters. This is why we will compare the root mean squared error with the mean absolute deviation as candidate focused measures. We have selected these two focused measures since they are designed to emphasize avoiding different types of errors; root mean square error places a much greater emphasis on avoiding large errors than mean absolute deviation. We do not consider relative measures such as mean absolute percentage error since the demand time series values do not change much from one portion of the simulation to the next.

Note that all forecast errors compare the forecast point estimate with the "actual" data value. The "actual" data value is obtained from the Poisson intensities and the appropriate cancellation and no-show probabilities from Baker and Collier. [18] Thus, all forecast errors involve comparing true net demand predictions with "actual" true net demand. So, "actual" true net demand is the sum of the Poisson intensities across all reading periods times (1 – cancellation rate)**(number of reading periods) times (1 – no-show probability).

Also, note that the purpose of this study is to generate insights, not absolute numbers for later use. Since the specific simulation parameter values from one hotel property to the next will vary from our mainline numbers, their best values to use for frozen periods, number of data points, internal measure of merit, and horizon length will likely differ from ours. What should generalize from this study, though, are the qualitative patterns in the optimal frozen periods-number of data points-measure of merit-horizon length used. For instance, it may be that the base smoothing constant in the Holt-Winter's forecaster is far more important to control tightly than the trend or seasonal constants. Likewise, it might be more important to freeze the Holt-Winter's smoothing constants longer than the booking profile percentages. Unlike the smoothing constant updates, where there is no a priori (from the literature or first principles) reason to hypothesize any outcome, one would expect a longer optimal frozen period for the Holt-Winter's forecaster than the booking profile forecaster. Since the booking profile forecaster is designed to be more reactive to recent trends. But for the most part, there are no a priori expectations about the results of this study. Below in the *Factor Description* section, we will discuss any a priori expectations about the results. Therefore, it is exploratory in nature; we will discern patterns in the results and then generate tentative explanations for them. That can be tested more formally in future research.

We have identified 16 of these data volume/frozen interval length/measure-of-merit/horizon length factors that could impact the revenue-generating performance of a revenue management system. From the baseline forecasting system that is standard in practice today (see Duffy), these 16 comprise all of the possible variables that could be changed to alter a parameter/frozen interval/measure-of-merit/horizon length in the system. [5] We will run a 16 factor Plackett-Burman screening design to isolate the most influential factors on revenue generating performance. Once the final regression equation is formed, we will explore it to generate insights about ideal forecaster settings. All main effects and two-way interactions will be evaluated via this screen. 30 independent replications will be performed for each point in the design, and the final uniform numbers in the simulation random number stream from one design point will be used as the starting seeds for the next design point to ensure independence.

3.1. Factor Descriptions

Here is the list of all 16 forecaster system variables that the revenue management system manager has control over to influence forecast errors directly, and thereby net revenue indirectly. In order to fit the size constraints of the *QuikSigma* software used to do the analysis, we have eliminated two of these factors from consideration based on judgment over relative impact on net revenue. [19] Thus, the factors with letters A-N are the ones that we will use for analysis. These names then translate to the initial two sets of Plackett-Burman runs detailed in Tables 1-2. The low values for all factors correspond to one week. Given that the seasonal cycle is one week on these data, the minimum amount of time between specific (*rda* combination) forecast parameter updates is one week. Therefore, our minimum factor level for all factor updates is one week. Since the number of reading periods equals two weeks of data, our maximum factor level must exceed that. So, we have chosen five weeks in general for our maximum factor level

Factor 1 (A): How many arrival build-ups (i.e., complete reading period sets) to use in estimating true net demands for the booking profile per rate class-duration-arrival day-of-week (rda) combination. Since opening/closing *rda*s for sale only happens once per reading period, it is best to use multiple build-ups to deal with the case where an *rda* was closed for sale during a reading period—the average observed demand when the *rda* was open across the buildups is then used as the true demand estimate for the reading period when the *rda* was closed. Note that if the *rda* is closed during all build-ups, then true net demand is estimated to be 0 for that *rda*. We use one arrival for the low value, five arrivals for the high value.

Factor 2: For how many future arrival build-ups are the averages in Factor 1 held constant before being updated again. We use every week for the low value, and every five weeks for the high value.

Factor 3 (B): Booking profile data points. Each reading period in a booking profile corresponds to an estimate of the cumulative fraction of total true net demand that typically occurs by the end of that reading period. The issue is how many data points to use in computing that fraction. Our low value is one data point, and our high value is five.

Factor 4 (C): Booking profile frozen interval. How often should the booking profile be updated, by rate class-duration-arrival day-of-week? We use once a week as the low value, and once every five weeks as the high value.

Factor 5 (D): Holt-Winter's base smoothing constant data points. When optimizing the base smoothing constant via enumeration with regards to either the root mean square error or mean absolute deviation, how many historical forecasts are used? The low value is one, and the high is five.

Factor 6 (E): Holt-Winter's base smoothing constant frozen interval. How often is this smoothing constant updated, per rate class-duration-arrival day-of-week combination? The low value is every week, and the high value is every five weeks.

Factor 7 (F): Holt-Winter's trend smoothing constant data points. When optimizing the trend smoothing constant via enumeration with regards to either the root mean square error or mean absolute deviation, how many historical forecasts are used? The low value is one, and the high is five.

Factor 8 (G): Holt-Winter's trend smoothing constant frozen interval. How often is this smoothing constant updated, per rate class-duration-arrival day-of-week combination? The low value is every week, and the high value is every five weeks.

Factor 9 (H): Holt-Winter's seasonal smoothing constant data points. When optimizing the seasonal smoothing constant via enumeration with regards to either the root mean square error or mean absolute deviation, how many historical forecasts are used? The low value is one, and the high is five.

Factor 10 (I): Holt-Winter's seasonal smoothing constant frozen interval. How often is this smoothing constant updated, per rate class-duration-arrival day-of-week combination? The low value is every week, and the high value is every five weeks.

Factor 11 (J): Forecast error measure. The low value here means that the root mean square error (RMSE) is used, high value is mean absolute deviation (MAD).

Factor 12 (K): Final forecast booking profile weight. The weight between 0 and 1 to be put on the booking profile forecast is

$$\frac{\frac{1}{booking\ profile\ forecast\ error}}{\frac{1}{booking\ profile\ forecasrt\ error}+\frac{1}{Holt-Winter's forecast\ error}}$$

These forecast errors are obtained from either the root mean square error or the mean absolute deviation. These errors are computed for each rate class-duration-arrival day-of-week combination. The question is how many forecasts to use in this computation? The low value is one, and the high value is five.

Factor 13 (L): Final forecast frozen interval. Once a forecast booking profile weight has been recomputed, how many weeks does it remain constant? The low value is one week, and the high value is five weeks.

Factor 14 (M): Final forecast error coefficient of variation data points. For each rate class-duration-arrival day-of-week, how many true net demand final forecasts are compared with the "actual" true net demands? One forecast is the low value, and five forecasts the high value.

Factor 15: Final forecast error coefficient of variation frozen interval. Once a coefficient of variation has been recomputed, how many weeks must elapse before the next recomputation? The low value is one week, and the high value is five weeks.

Factor 16 (N): Forecast coverage length. How many weeks should each forecast—booking profile, Holt-Winter's, and final—go out into the future? The low value is one week, and the high value is two weeks. The high value is two weeks since the number of reading periods covers two weeks.

3.2. Data Analysis

For each of the 14 design points in Table 1, we will run 30 replications on the simulation to generate the regression dependent variable value as the mean net revenue. Likewise for the 14 design points in Table 2. Note that per Plackett-Burman methodology the signs for each factor are reversed between Table 1 and Table 2. This is done to generate a contrast for each of the 14 factors. Thus, once the mean revenues for those 32 design points has been generated, we will examine the *QuikSigma* output to determine (a) which main effects and two-way interactions had regression coefficient signs that switched from Table 1 to Table 2, (b) which main effects and two-way interactions had statistically insignificant coefficient signs on either Tables 1 or 2, and (c) which main effect and interactions had the wrong signs from either Tables 1 and 2. [19] Only those main effects-interactions whose signs remain consistent, correct, and significant across both Tables 1 and 2 will be considered further. Moreover, any confounding between coefficients will be resolved with further design points using judgment to determine the factor effect that underlies the coefficient.

Once all of these steps have been performed, a final regression using all of the data generated thus far on the surviving main effects and two-way interactions will be performed. We will then use this final regression equation to do exploratory analysis on, so that insights about optimal forecast system design can be generated.

We will repeat these steps for a 2^2 full factorial experiment, where the two factors are the demand intensity from Pimentel *et al.* and a gap measure from a concatenation of [18]. A low gap represents both peak gap and revenue gap at their low levels, and a high gap represents both of these factors at their high levels. We decided to collapse peak gap and revenue gap into one dimension since one replication of the simulation takes about 12 hours of computing time on the following computer: Intel Corel 15-4590 CPU @3.30 GHZ processor; 8.00 GB of installed RAM; 64-bit operating system x64-based processor. Since peak gap and revenue gap tend to move together in practice according to Duffy these four design points cover the complete range of operating conditions that a property might face. [5]

4. Results

4.1. High Demand Intensity-High Gap Case

The first round of the Plackett-Burman screen revealed only one factor that is statistically significant in predicting nightly net revenue: a factor where A * C, D * F, E * G, H * J, I * K, L * N, and M * C are all confounded with one another. Round 2 also reveals that this factor with these confounded elements is the only significant factor. For

Round 3, we decided to replicate the Round 2 design with factor L's sign flipped so that L * N would not be confounded with any of these other interactions. The rationale for isolating on L * N in Round 3: a process of elimination. A * C represents the interaction of how often the observed bookings to true demand conversion is performed with the length of the frozen interval on the booking profile forecast. Since these two factors are unrelated, this interaction has no interpretation. D * F represents the interaction of the amount of data to use in updating the Holt-Winter's base smoothing constant with the amount of data to use in updating the Holt-Winter's trend smoothing constant. There is no practical interpretation of this interaction. E * G represents the interaction between the length of the frozen period for updating the Holt-Winter's base smoothing constant with the corresponding trend frozen period length. No interpretation here. H * J represents the interaction between the amount of data used in the Holt-Winter's seasonal smoothing constant and whether to use RMSE or MAD as the forecast error measure. Again, these two factors are unrelated, so this interaction has no interpretation. I * K represents the interaction between the Holt-Winter's frozen interval length for updating the seasonal factor and the amount of data to use in optimizing the final forecast weight to put on the Holt-Winter's method relative to the booking profile method. No interpretation of I * K makes sense. Finally, M * C represents the interaction between amount of data used in the update of the coefficient of variation for the final forecast and the length of the frozen interval for the booking profile forecast. No interpretation is possible.

Thus, the only remaining interaction is L * N. This interaction between the length of the final forecast frozen interval and the final forecast planning horizon. This interaction should be negative since shorter horizons are more accurate and stable, so freezing that short horizon forecast makes the overall forecasting scheme less reactive to noise. However, longer forecasting horizons mean more error overall, so there is more of a need to update these forecasts more often.

The result of Round 3 was that the regression coefficient signs on A * C, D * F, E * G, H * J, I * K, and M * C all switched from Round 2. This means that these interactions are unreliable as true predictors of nightly net revenue, so that leaves us with L * N as the lone factor that is significant in predicting net revenue.

The adjusted R^2 is 24.45% when one concatenates the data from Rounds 1-3 and regresses average nightly net revenue against L * N.

4.2. High Demand Intensity-Low Gap Case

Rounds 1-3 progressed in identical fashion to the High Intensity-High Gap case in terms of significant factors and the set of confounded interactions that comprise each factor. Thus, at the end of Round 3, only the L * N interaction was left as meaningful. The concatenation of Rounds 1-3 for the final regression yielded both L and L * N as significant factors, with an adjusted R^2 of 47.37%. The sign on L * N was negative as in the High Demand-High Gap case. The sign on L is positive, indicating that longer frozen periods for the final forecast are more desirable in general.

4.3. Low Demand Intensity-High Gap Case

Rounds 1-3 proceeded in identical fashion as the prior two cases in terms of significant effects and confounding interactions at each Round. The regression model with Rounds 1-3 concatenated yielded an R^2 of 30.7% with only the L * N term significant. And with the negative coefficient.

4.4. Low Demand Intensity–Low Gap Case

Rounds 1-3 proceeded as in the other three scenarios. The final regression model had an R^2 of 14.95% with only the L * N term significant, and with a negative coefficient.

5. Discussion

We utilize a simulation model of a hotel reservation system, validated by Duffy, that has all of the functionality of a real world revenue management system for transient customers. [5] The size of the simulation is proportionately scaled down to permit computational experiments in a reasonable timeframe, without compromising the full richness of interactions between all system components.

The demand forecasting component of a hotel revenue management system has received scant attention in the literature. What little that has been published focuses on specific techniques of forecasting, rather than how the forecasting system as a whole should be parameterized. This study fill this void by exhaustively compiling and analyzing all of the parameters of a current real world forecasting system, with the objective of separating the critical few parameters from the not-so-critical many. And then obtaining some insights about forecaster parameterization.

We executed a 2^2 full factorial design, where the two factors are the level of capacity saturation and the degree to which market segments differ in terms of price points and booking patterns. Thus, the full range of operating environments in a hotel are considered, according to Duffy. [5]

Within each design point, we ran a Plackett-Burman screening experiment on the list of 14 forecasting system parameters that effectively completely define the operations of such a system. The screen shows that only two parameters really influence the revenue generation performance of the system: how long the final combined forecast is frozen for, and the final forecaster planning horizon length. These interact strongly to influence net revenue generation. The negative interaction is due to the fact that shorter forecasting horizons lead inherently to more accurate and stable forecasts, so freezing these forecasts for more time makes sense—there is no need to update the forecasts more frequently to compensate for large errors. Whereas the longer forecasting horizons yield more error, so a shorter frozen period will enable faster corrections to these errors.

Thus, we have essentially prioritized efforts on improving forecasting systems. The other parameters besides final forecast horizon length and frozen period do not need to be scrutinized; they can be left at default settings. These parameters include the amount of historical data to use for (1) all three Holt-Winter's smoothing constants, (2) the booking profile, (3) the estimation of true demand from censored bookings, (4) the short-term versus long-term forecaster weights, and (5) the construction of forecast error coefficient of variation. Also, how long various forecasts and their associated parameters are frozen for do not need to be considered: (1) all Holt-Winter's smoothing constants, (2) all booking profile forecasts, (3) all Holt-Winter's forecasts, and (4) the forecast error measures.

6. Conclusion

We have utilized a hotel reservation system simulation as our method for estimating nightly net revenue, as a function of assorted forecasting system characteristics. This simulation was validated by a major hotel chain. For each of four combinations of hotel capacity saturation and degree of market segment differentiation, we ran a screening design on 16 forecast system parameters. Our results indicate that the vast majority of parameter settings can be ignored in practice, so long as they stay within the broad range of reasonable settings detailed in our analysis. Thus, (a) the booking profile (i.e., short-term, reactive forecaster component) settings—how much data to use in updating, how often to update—are irrelevant, (b) all three parameters of the long-term forecaster—base term, trend term, seasonal term—can have both their volume of data for updating and frequency of updating settings ignored, (c) which measure of merit in evaluating the relative importance of long-term versus short-term forecasting efficacy for the purpose of combining them into a final forecast—the choice of measure is immaterial, (d) the amount of data used to weigh the relative importance of long versus short-term forecaster importance in the final forecast is irrelevant, (e) the amount of data and frequency of update of the forecast error coefficient of variation (for eventual input into the room allocation optimizer) is immaterial.

Therefore, this analysis greatly simplifies the focus the analyst must have in tuning the forecasting system. They only have to be concerned with conditions involving the final forecast (that combines the long-term and short-term forecasts): the length of the final forecast horizon, and how long the corresponding this final forecaster is frozen for.

Although the current state-of-the-art generally involves relatively simple techniques--the Holt-Winter's three parameter exponential smoothing model for long-term forecasts, and the booking profile forecast for short-term forecasts—exploration of more advanced techniques is warranted for future research. These simple techniques often work well because they do not exacerbate forecast errors due to overfitting to noise. A good topic for future research would be to see under what specific conditions overfitting is not such a problem, so that more advanced techniques such as causal methods and artificial intelligence based approaches might work better.

Appendix

Table 1. *Factor Levels for first 16 Plackett-Burman screening design runs.*

(-1 denotes the factor at its low level, 1 at its high level)

A	B	C	D	E	F	G	H	I	J	K	L	M N	
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
-1	-1	-1	-1	-1	-1	-1	1	1	1	1	1	1	1
-1	-1	-1	1	1	1	1	1	1	1	1	-1	-1	-1
-1	-1	-1	1	1	1	1	-1	-1	-1	-1	1	1	1
-1	1	1	1	1	-1	-1	-1	-1	1	1	1	1	-1
-1	1	1	1	1	-1	-1	1	1	-1	-1	-1	-1	1
-1	1	1	-1	-1	1	1	1	1	-1	-1	1	1	-1
-1	1	1	-1	-1	1	1	-1	-1	1	1	-1	-1	1
1	1	-1	-1	1	1	-1	-1	1	1	-1	-1	1	1
1	1	-1	-1	1	1	-1	1	-1	-1	1	1	-1	-1
1	1	-1	1	-1	-1	1	1	-1	-1	1	-1	1	1
1	1	-1	1	-1	-1	1	-1	1	1	-1	1	-1	-1
1	-1	1	1	-1	1	-1	-1	1	-1	1	1	-1	1
1	-1	1	1	-1	1	-1	1	-1	1	-1	-1	1	-1
1	-1	1	-1	1	-1	1	1	-1	1	-1	1	-1	1
1	-1	1	-1	1	-1	1	-1	1	-1	1	-1	1	-1

Table 2. *Factor Levels for second 16 Plackett-Burman screening design runs.*

(-1 denotes the factor at its low level, 1 at its high level)

A	B	C	D	E	F	G	H	I	J	K	L	M N		
1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
1	1	1	1	1	1	1	-1	-1	-1	-1	-1	-1	-1	-1
1	1	1	-1	-1	-1	-1	-1	-1	-1	-1	1	1	1	1
1	1	1	-1	-1	-1	-1	1	1	1	1	-1	-1	-1	-1
1	-1	-1	-1	-1	1	1	1	1	-1	-1	-1	-1	1	1
1	-1	-1	-1	-1	1	1	-1	-1	1	1	1	1	-1	-1
1	-1	-1	1	1	-1	-1	-1	-1	1	1	-1	-1	1	1
1	-1	-1	1	1	-1	-1	1	1	-1	-1	1	1	-1	-1
-1	-1	1	1	-1	-1	1	1	-1	-1	1	1	-1	-1	1
-1	-1	1	1	-1	-1	1	-1	1	1	-1	-1	1	1	-1
-1	-1	1	-1	1	1	-1	-1	1	1	-1	1	-1	-1	1
-1	-1	1	-1	1	1	-1	1	-1	-1	1	-1	1	1	-1
-1	1	-1	-1	1	-1	1	1	-1	1	-1	-1	1	-1	1
-1	1	-1	-1	1	-1	1	-1	1	-1	1	1	-1	1	-1
-1	1	-1	1	-1	1	-1	-1	1	-1	1	-1	1	-1	1
-1	1	-1	1	-1	1	-1	1	-1	1	-1	1	-1	1	-1

References

[1] Kunnumkal, S., Talluri, K. (2019). A strong Lagrangian relaxation for general discrete-choice network revenue management. *Computational Optimization & Applications*, 73 (1), 275-310.

[2] Zhang, D., Weatherford, L. (2016). Dynamic pricing for network revenue management: A new approach and application in the hotel industry. *INFORMS Journal on Computing*, 29 (1), 18-35.

[3] Druckman, A. (2003). Rent check. Journal of Property Management, 68 (1), 51.

[4] Vives, A., Jacob, M., Aguilo, E. (2019). Online hotel demand model and own-price elasticities: an empirical application in a mature resort destination. *Tourism Economics*, 25 (5), 670-694.

[5] Duffy (2015). Personal communication with Marriott's corporate revenue management group (June 15, Bethesda, MD, USA). Kathleen Duffy is the contact.

[6] Lai, K., Ng, W. (2005). A stochastic approach to hotel revenue optimization. *Computers and* Operations Research, 32, 1059-1072.

[7] Weatherford, L., Kimes, S. (2003). A comparison of forecasting methods for hotel revenue management. International Journal of Forecasting, 19 (3), 401-415.

[8] Weatherford, L. & Belobaba, P. (2002). Revenue impacts of fare input and demand forecast accuracy in airline yield management. Journal of the Operational Research Society, 53 (8), 811-821.

[9] Weatherford, L, Kimes, S., Scott, D. (2001). Forecasting for hotel revenue management. Cornell Hotel and Restaurant Administration Quarterly, 42 (4), 53-64.

[10] Azadeh, S., Marcotte, P., Savard, G. (2014). A taxonomy of demand uncensoring methods is revenue management. Journal of Revenue and Pricing Management, 13, 1-17.

[11] van Ryzin, G., Vulcano, G. (2017). Technical note—an expectation-maximization method to estimate a rank-based choice model of demand. *Operations Research*, 65 (2), 396-407.

[12] Yuksel, S. (2007). An integrated forecasting approach to hotel demand. Mathematical and Computer Modeling, 46 (7), 1063-1070.

[13] Huang, Y., Ge, Y., Zhang, X., Xu, Y. (2013). Overbooking for parallel flights with transference. *International Journal of Production Economics*, 144 (2), 582-589.

[14] Phumchurri, N., Maneesophon, P. (2014). Optimal overbooking decisions for hotel room revenue management. Journal of Hospitality and Tourism Technology, 5 (3), 261-277.

[15] Fouad, A. Atiya, A., Bayoumi, A., Saleh, M. (2014). A simulation-based overbooking approach to hotel revenue management. Institute of Electrical and Electronic Engineers, Dec 29, 61-69.

[16] Koch, S. (2017). Least squares approximate policy iteration for learning bid prices in choice-based revenue management. *Computers & Operations Research*, 77, 240-253.

[17] Pimentel, V., Aizezikali, A., Baker T. (2016). The Impact of Simultaneous Overbooking-Allocation on Explicit Price Setting Methods in Hotel Revenue Management. Working paper, Washington State University, Pullman, WA.

[18] Baker, T., Collier, D. (2003). The benefits of optimizing prices to manage demand in hotel revenue management systems. Production and Operations Management, 12 (4), 502-518.

[19] *QuikSigma* (2011). Software for Six Sigma project management. Promontory Management Group (Layton, UT, USA).