II International Conference on Industrial Engineering and Industrial Management
XII Congreso de Ingeniería de Organización
September 3-5, 2008, Burgos, Spain

Demand planning in hotel management Decision Support Systems

Alfonso Durán Heras[1], Isabel García Gutiérrez[1]

[1] Ingeniería de Organización. Escuela Politécnica Superior. Univ. Carlos III de Madrid. Avda. Universidad, 30. 28911 Leganés, Madrid. alfonso.duran@uc3m.es, isabel.garcia@uc3m.es

Abstract

Yield Management based Decision Support Systems offer significant potential for small and medium hotels and hotel chains in terms of better capacity utilization, profitability and ability to compete with larger, international chains. However, the very specific and exacting demand planning requirements of these systems hinder their application by SME. This paper analyzes the differential pricing theoretical foundations of Yield Management to derive the challenging demand planning requirements it imposes. The major stumbling blocks found in meeting these requirements are then studied. The paper then proposes a design approach for the demand planning modules of such a DSS, developed in the framework of both a larger YM based DSS research project and a cooperative project with a small hotel chain, aimed at surmounting these obstacles. This design might provide a viable step towards the actual implementation of these DSS by the Spanish hotel SME.

Keywords: Demand Planning. Yield management. Decision Support Systems.

1. Differential pricing and demand forecasting in Yield Management based DSS

Facing mounting competitive pressures, hotels are gradually adopting managerial techniques aimed at extracting the maximum possible revenue out of resources (hotels / hotel rooms) whose costs are basically fixed regardless of their degree of utilization. Thus, their Decision Support Systems (DSS) increasingly rely on Yield Management (YM) techniques, also known as Revenue Management (RM) (Baker and Collier, 2003; Chiang et al, 2007). YM involves dynamic methods and optimization heuristics to plan demand, allocate perishable assets (for hotels, the hotel rooms) across rate classes, decide when and by how much to overbook and what price to charge different rate classes in order to maximize revenues for the firm. YM algorithms form the core of the Decision Support Systems (DSS) routinely utilized by companies such as US Airlines and Delta Airlines, which reportedly increased revenue by US$500 and $300 million, respectively (Boyd, 1998) or Marriott Hotel, that gained US$100 million additional annual revenue (Elliott, 2003).

The key distinction between YM/RM and traditional pricing approaches is its focus on segmented, differentiated pricing. YM focuses on identifying the different willingness to pay of various customer segments, then tries to extract the maximum possible revenue by charging each customer / customer category a price as close as possible to the value for the consumer of that product / service. Its theoretical foundation lies in the microeconomic differential pricing curve. Figure 1 shows how, for a given available output q0, such as a certain number of rooms available for a particular night, revenue without price discrimination equals q0 * p0, p0 being the price at which the market demand equals q0. (If the demand is inelastic, that revenue can be increased by charging a higher price, as a monopoly would. On the other hand, if the producer mistakenly charges a price lower than p0 , or the demand is elastic and he charges more than p0, the revenue will be lower than q0 * p0). The difference between p0 and the demand curve shows the difference between what these consumers would have been willing to pay and what they actually paid; it is called customer surplus.

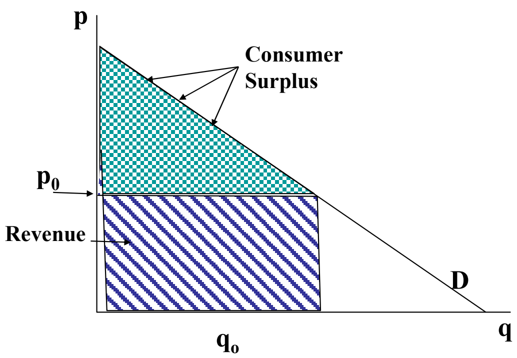

Figure 1. Revenue and consumer / customer surplus without price discrimination

Differential pricing is an attempt by the producer to appropriate at least part of that customer surplus. Figure 2 shows how even an imperfect price discrimination allows producers to appropriate a significant share of the customer surplus.

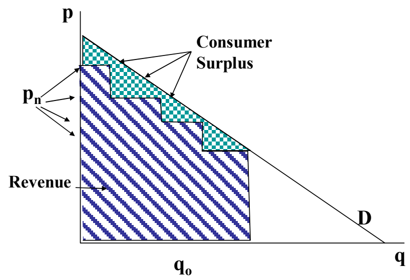

Figure 2. Revenue and consumer / customer surplus with price discrimination

Appropriating the customer surplus is a zero sum game. However, as mentioned before, given inelastic demand and non-differentiated pricing, it is in the producer's interest to sell less than the available number or rooms in order to secure a higher price for each of them (thus, q0 in Figure 1 would be less than the number of available rooms). On the other hand, if the producer has the ability to effectively discriminate in prices, and costs are fixed, the producer has an incentive to increase the number of rooms actually sold. This moves part of the non-shaded right-hand triangle into a combination of Revenue and Costumer Surplus. That is a positive-sum game, and in some cases it can be shown to result in both higher Revenue and higher Customer Surplus than without applying YM (for example, in cases where the maximum achievable revenue without applying price discrimination does not cover the fixed costs) (Varian, 1996 and 2007)

Pigou's taxonomy describes three degrees of price discrimination: First-degree (price might differ from person to person, and also among the various units sold to the same person – in its extreme form it leads to perfect discrimination), second-degree (price might be different depending on the number of units bought by the same person, but there is no discrimination among individuals – e.g. volume discounts) and third-degree (price might differ from person to person, but is independent of the number of units bought, e.g. student discounts).

In real life, hotel YM/RM systems can not achieve perfect discrimination, whereby the supply curve would be coincident with the demand curve, and the whole customer surplus would be appropriated by the producer. Even in the simplified case in which each customer only buys one or zero units, that would require: a) Setting a different price for each customer b)figuring out her specific maximum price ("willingness to pay"), represented in the demand curve by the precise point, or q increment, accounted for by this customer c) ensuring perfect segmentation / isolation among segments, i.e. ensuring she does not gain access to the product or service at the lower rate which others are enjoying.

A more realistic target, as depicted in Figure 2, is a stepwise differential pricing approach, whereby customers are grouped in customer segments that can somehow be at least partially isolated, and then each customer segment is charged a different price according to the willingness to pay of its constituents. This approach requires identifying segmentation criteria (several segmentation criteria can be used simultaneously) that meet at least some of these conditions:

- Clusters together customers with similar willingness to pay; the ideal criterion would group customers represented by contiguous points in the Demand curve.

- The price that customers in a given segment are willing to pay can be established.

- Segments can be at least partially isolated, thus preventing a customer in a segment to which a higher price is offered from obtaining the lower prices being offered to other segments.

The third condition is the only prerequisite for the application of YM. It can be met by linking the customer segmentation to verifiable customer traits (age, profession...); the usefulness of these criteria depends on their correlation with the willingness to pay. It can also be achieved through distribution channel segmentation; however, since that approach used to rely on the consumer's unawareness of the existence of a cheaper channel, the increasing transparency brought about by Internet is hindering its application. The most common approach, however, is to intentionally hamper the attractiveness for the consumer of the product of service (even if the associated production cost is not decreased), in various degrees, to create an artificial "product range". A well-known example from the airline industry, where YM/RM was born, are tickets requiring a Saturday night stay. That does not directly benefit the airline, but it allows it to separate business travellers from holiday travellers. In the case of hotels, that is normally achieved through the advanced reservation lead time (also very commonly used by airlines): the longer the anticipation in the reservation, the lower the price. That can be implemented by actually adjusting the price, or by using a quota of cheaper rates (through vouchers...) that sells out before full rate rooms do.

The first and second conditions are the main focus of this paper. The discussion above shows that the ability of an YM/RM based DSS to maximize revenue is contingent on the accuracy of the predictions of the patterns and features (e.g. price sensitivity and its relation with the segmentation criteria) of the demand; thus firms using these DSS must exploit every relevant source of information to increase this accuracy. The increasingly utilized term demand planning, also sometimes referred to as supply chain forecasting, encompasses both statistical forecasting and judgemental methods (incorporating intuitive judgements, opinions and probability estimates). Therefore, the "Advanced DSS for hotel management" multi-year, multi-centre research project approved in 2005 by the Spanish Ministry of Education, within which this study is conducted, contemplates demand planning as one of the DSS's key building blocks.

2. Demand planning issues

Demand planning for YM based hotel DSS faces significant challenges. Different hotels use different business process designs for their room reservation process (e.g., keep price for each rate class constant or allow it to float, penalize cancellations or not...). The appropriate algorithm is contingent on that design, and, in turn, each YM algorithm has specific forecasting requirements. Furthermore, for the reasons discussed before, these forecasting requirements go well beyond conventional estimations of demand volumes. If reservation lead time is used as the segmentation criterion (in can be used in combination with other criteria), they require a forecast of the number of customers that, each day, will make an advanced booking for each future date, thus resulting in a bi-dimensional matrix. That matrix must be further broken down by the customer's demand class, subpopulation, or any other proxy for his readiness to pay, thus leading to a tri-dimensional forecast. The definition and details of each of these dimensions or axes is dependent on the specific algorithm.

This implies that the application of statistical forecasting in this environment must surmount significant hurdles. In initial stages of the DSS implementation, even in hotels that keep extensive historical data it is difficult to re-structure that in the required format. After implementation, this impediment can be gradually overcome as an ad-hoc history builds up with the required structure.

Application of judgemental methods is not without its own stumbling blocks. Besides its inherent subjectivity, difficulties include the complex structure and level of detail required in the forecasted data, its probabilistic nature and the fact that, in most YM algorithms, data for each "execution" date (i.e., for each night) is initially forecasted but must then be updated after each new reservation.

On the other hand, accuracy in the various stages of the demand planning process is critical for the effective application of an YM/RM DSS. At an initial design stage, when the segmentation criteria must be chosen, an inaccurate estimation of the correlation between price sensitivity and a candidate criterion might result in the wrong choice of criteria. Once a given criterion is chosen (e.g. reservation lead time), the YM algorithm will suggest a price for each "segment" (in this example, for each lead time period: 20 days anticipation, 15 days anticipation ...) based on the forecasted customer mix for that period and for each of the following ones. From the producer's perspective, that price will be too low for some customers in that segment (they will book the room at that price, but they would have been willing to book it also at a higher price) and might be too high for others (they would not book the room at that price, but they would have booked it at a lower price that was still worthwhile for the producer). The right choice requires striking a balance between those two sets of customers, and is obviously contingent on the accuracy of their forecast.

3. Demand planning process

The appropriate approach to tackle those challenges is contingent on the characteristics of the hotel, particularly its size. Since a primary objective of the abovementioned research project was to develop solutions aimed at the Small and Medium Enterprise (SME) range, a collaborative project has been started with a Madrid based small chain of hotels to identify a realistic design for the demand planning process required for an YM DSS. As a result of this collaboration, an ad-hoc demand planning process specifically aimed at these challenges has been designed (Figure 3). It includes an adjustable combination of system-aided judgemental forecasting with data-supported adjustment and updating, as described in the following sections.

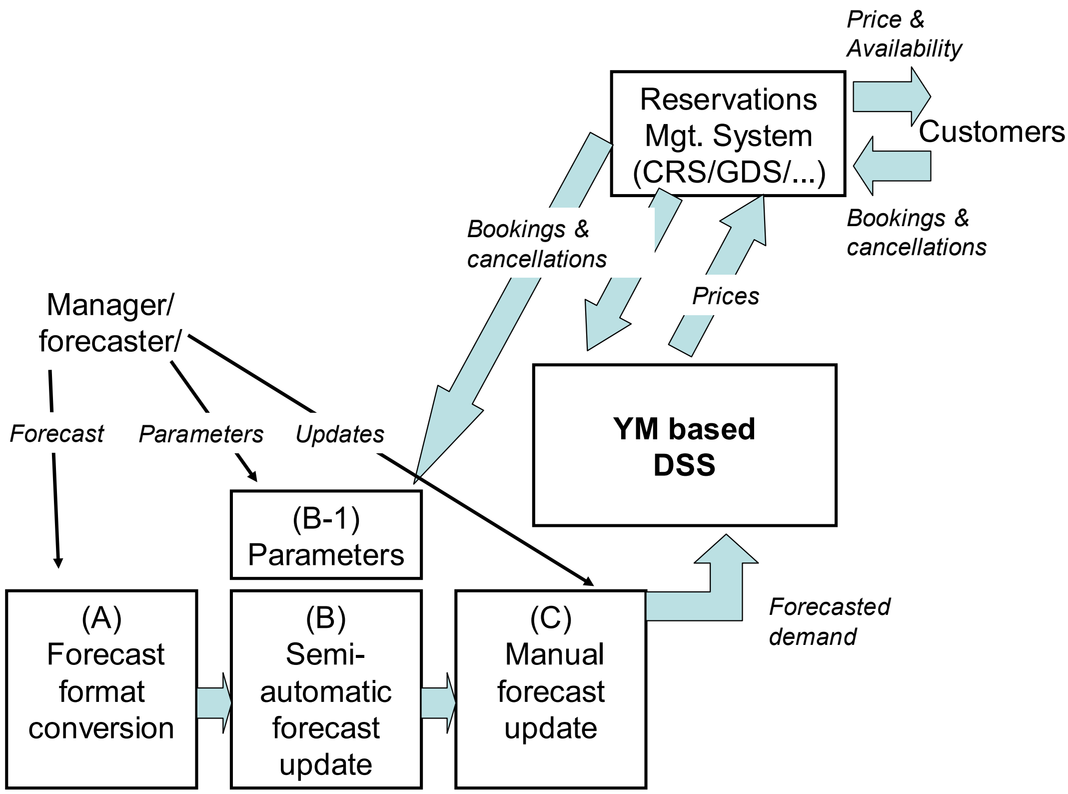

Figure 3. Demand planning process in YM-based DSS for hotel SME

3.1. Reservation management system and YM based DSS

Figure 3 shows an YM based DSS that sets prices (a human or software based agent, not depicted here, might filter/approve those DSS-suggested prices before actually implementing them). These prices are fed into the system actually handling reservation requests (computer reservations system (CRS) or global distribution systems (GDS) will nowadays normally handle reservations for airline tickets, hotel rooms and rental cars). This reservation management system handles the interaction with the consumers, provides them with price and availability information and processes their booking requests. In the case of the small chain of hotels analyzed in this project, that functionality is provided by a portal operating in Application Service Provider (ASP) mode, which is a very common approach among SME in the hotel sector. This reservation management system maintains up to date information on room bookings and availability, and provides it (online or otherwise) to: a) the YM based DSS and b) the demand planning modules.

As discussed above, different YM algorithms have different requirements in terms of inputs. However, algorithms potentially applicable in a SME in the hotel sector will most likely issue their price decision recommendations in two stages, as is the case for the family of algorithms initially selected for this project. Assuming reservations for a given night (the execution date) are only accepted a certain Lead Time (LT) before, the algorithm must first issue a recommendation regarding pricing decisions at the beginning of that reservation window (i.e., LT days before the night whose reservations it is analyzing, “execution date – LT”). Obviously, at that time, no reservations will exist, and therefore the only input at this stage is the demand forecast.

From a theoretical standpoint, the algorithm could treat this decision as final (i.e., not review it as the execution date approaches). However, no real life hotel owner implementing an YM based DSS is likely to forgo the profit potential stemming from knowing the evolution of actual bookings as the time advances. Thus, the DSS will normally periodically re-execute its YM algorithm, using the actual updated firm bookings info (and, consequently, availability), and using a forecast for the remaining of the lead time.

Therefore, as depicted in Figure 3, the YM based DSS will receive two inputs: demand forecast from the demand planning module (initial forecast and, potentially, updated forecasts as time advances) and updated bookings and availability information from the reservation management system.

3.2. Forecast and forecast format conversion

As discussed in the “Demand planning issues” section, extrapolating through quantitative methods the required forecast from the (generally scant) existing historical data in SME in the hotel sectors is laden with hurdles. An additional stumbling block in actually implementing that approach is the foreseeable reluctance by the hotel owner to relinquish control, not just to a “black box” YM algorithm, but also to one based on a demand forecast coming from a second “black box”. Therefore, the chosen design is based on feeding the system with a forecast provided by the owner/manager/ forecaster, at least until an ad-hoc history builds up over time with the required structure after implementation (at that point, the initial forecast discussed in this section could be “switched” in the design to a system-generated one). The people providing the inputs for this forecast are likely to be the same that previously took the pricing decisions (decision implicitly based in an estimation of demand).

However, the forecaster (in an SME, the forecaster will most likely be either the owner or the manager) can not be reasonably expected to fill up a forecast matrix that, given the peculiar forecasting requirements of YM algorithms, for a given type of room that accepts bookings 2 months in advance and contemplates 2 alternative prices, would require completing 3.800 forecasted values for each month. It is thus imperative that a format conversion module is included in the design, allowing the forecaster to introduce a limited amount of aggregated data that can then be broken down to the required detail level.

A persistent problem while trying to design a generalized but implementable YM based solution is the fact that so many aspects are contingent on the specific YM algorithm used. In the case of this module, this general problem boils down to the different forecasting requirements posed by different algorithms, and more specifically to how Willingness To Pay (WTP) is forecasted. For the purpose of this project, the option to discretely (as opposed to continuously) measuring WTP by breaking down the demand for a given time period in “W” WTP brackets is reckoned to be a sufficiently generalized approach. This would imply forecasting, for each room type, for each night (execution date), and for each “anticipation” (the subdivisions into which the maximum booking anticipation or “Lead Time” is split; it will generally be each of the days of the LT), “W” values, each representing the number of customers prepared to make a booking for that room type and that night with that anticipation and for a price in that price bracket.

Once the structure of the desired output (the forecast matrix) has been decided, the design objective of the forecast format conversion can be expressed as: Identifying a relatively small set of data elements that the forecaster can estimate with reasonable accuracy and without excessive effort, and that can be processed to create the desired output format in a way that utilizes to the maximum possible extent the knowledge of the forecaster.

Two basic approaches can be combined to achieve that objective: factor decomposition and comparative analysis of the time patterns for different factors.

Factor decomposition refers to establishing several "influencing factors" that combined in a certain manner (multiplicative, additive ....) explain a significant part of the values of a variable. Under certain circumstances, techniques such as Principal Component Analysis (PCA) can be applied to derive from correlated components a new set of uncorrelated (i.e., orthogonal) factors, sometimes called principal components or eigenvectors.

Its relevance in this case is that if the variable to be forecasted can be approximately broken down into several multiplicative factors, each of which has an intuitive or business meaning for the forecaster, it might be feasible to ask the forecaster to estimate these factors, and from them to deduce a much larger number of forecasted value points (a "combinatorial implosion"). Specifically, since the forecasting problem being tackled in this module (the large number of forecasted value point required) precisely stems from the "curse of dimensionality" (the forecast matrix includes three dimension: execution date, lead time and willingness to pay), a promising approach is to look for factors that can be deemed approximately independent of some of these dimensions. It should be noted that this approach is very different from that taken in normal multivariate statistics, in which several value points, one from each independent variable, are used to estimate a single value point of a dependent variable.

The "comparative analysis of the time patterns for different factors" is precisely an attempt to choose some of the factors so that they are independent of the "execution date", so that fewer value points need to be forecasted. As an example, if we estimate that the distribution pattern for the reservations with respect to the booking anticipation (i.e., which % of the bookings takes place 15 days before the pernoctation date, which % is booked 10 days in advance, ...) can be considered independent of the specific pernoctation or execution date, then that distribution can be estimated just once and repetitively applied for each pernocation date.

The initial design chosen for this project contemplates a purely multiplicative model which does not explicitly contemplate correlations among the factors. The rationale for that approach is the difficulty for the forecaster of initially estimating these correlations. As the history associated with the application of this approach builds up, it will be possible to complement the model with quantitative, data-based estimations of the correlations.

As discussed above, for each hotel and room type, the variable to be forecasted is the number "rd,l,w" of customers willing to make a reservation.

The three dimensions are:

- d, the execution date (night for which the reservation is being made). It is modelled as a moving window, measured from the current date. Its range of values would be determined by the hotel's planning horizon, D days. However, in order to apply YM that planning horizon can not be shorter than the advanced reservation lead time (LT, maximum number of days before the execution date that a room can be booked). On the other hand, from an YM perspective a planning horizon longer than LT is irrelevant, thus it can be taken that $1 \leq d \leq LT$.

- l, anticipation lead time, number of days between the booking date and the execution date. $0 \leq l \leq (LT-1)$.

- w, Willingness To Pay bracket. As discussed, the price range is broken down into W brackets, associated with W prices. The first (top) bracket refers to customers willing to make a booking at the first (maximum) price. Each subsequent bracket refers to customers willing to make a booking at that price, but not at the next higher price. $1 \leq w \leq W$.

As an example to illustrate the proposed design, for a given hotel and a given type of room that accepts bookings 2 months in advance and measures willingness to pay through 3 different price brackets, the forecast moving window would include 10.800 forecasted values. (LT*LT*W, where LT=60 and W=3).

The proposed model introduces an additional, instrumental dimension: c, or customer category. The forecaster is asked to choose a customer classification criterion that meets one of the previously stated conditions for segmentation criteria (i.e. clusters together customers with similar willingness to pay) but does not meet the indispensable one (segments can be at least partially isolated, price-wise), and therefore can not be used directly as an YM segmentation criterion. Besides being highly correlated with the willingness to pay, this classification criterion must have a clear intuitive or business meaning for the forecaster, so that different values of the factors can be estimated for the C subpopulations or categories into which this criterion classifies the customers. In many hotels this criterion can be the distinction between business guests and tourism guests (thus C=2), however in some hotels the appropriate criterion might be a different one (e.g. national vs. foreigner, ...).

The forecaster is asked to estimate:

- For each night (execution or pernoctation date), the number nd,c of customers from each customer category that would book a room if the price is set at the lowest rate. (LT*C values).

- For each customer category, the distribution pattern for the reservations with respect to the booking anticipation for a "generic" execution date, i.e. the % of the total bookings for that execution date that would take place "l" days before d, pc,l (C*LT values).

- For each customer category, the percentage (out of those that would book a room if the price is set at the lowest rate) that would fall in each willingness-to-pay bracket, qw,c (thus assuming it is not dependent on either d or l) (W*C values)

The procedure to complete the forecast matrix would then compute:

$rd,l,w = \Sigma$ (from c=1 to c=C) $nd,c^* \, pc,l * qw,c$

In the example cited above, if two customer categories are specified (business and tourist guests, thus C=2), the 10.800 forecasted rd,l,w values would be derived from 246 estimated values. From there onwards, only 60 additional values would need to be estimated for each subsequent month.

(LT*C + C*LT+ W*C values = 120 + 120 + 6 since LT=60, W=3 and C=2) (subsequent months: 30*C).

This is the simplest model, applicable at the initial stages and in the smaller hotels. It can be further developed at a reasonable additional operational cost by:

- Assuming that the percentage in each willingness-to-pay bracket, q, is also dependent on the anticipation, l. (Thus, qw,c would become qw,c,l and, instead of W*C values , it would take W*C*LT values).
- Rather than assuming that the distribution pattern for the reservations with respect to the booking anticipation, pc,l, is independent from the execution date, create S representative distributions for each customer category (for example, one for weekends, one for weekdays and a third one for holiday periods, thus S=3). For each execution date, the forecaster must specify an sd value, thus indicating which distribution is applicable on that date (Thus, pc,l would become pc,l,s, sd would also be estimated, and, instead C*LT values, the forecaster would need to estimate C*LT*S + LT values)

The procedure to complete the forecast matrix would then compute:

$rd,l,w = \Sigma$ (from c=1 to c=C) $nd,c^* \, pc,l,s * qw,c,l$, where, for each d, s is set to sd.

In the previous example, if both developments are included, assuming S=3 the total number of values to be estimated for the forecast window would be 900; from there onwards, only 90 additional values would need to be estimated for each subsequent month.

(LT*C + C*LT*S + LT + W*C*LT values = 120 + 360 + 60 + 360, since LT=60, W=3, C=2 and S=3) (subsequent months: 30*C + 30)

3.3. Semi-automatic forecast update and manual forecast update

Equivalent designs have been created for the semi-automatic forecast update and the manual forecast update modules, whose details can not be explained in this paper due to space limitations.

The basic principle behind the semi-automatic forecast update is, once real booking data for a given execution date becomes gradually available (because that execution date is less than LT days away), to update the remaining rd,l,w based on the divergence between actual bookings in the already elapsed portion of the LT and the forecasted bookings for that period. These forecasted bookings would be the sum of rd,l,w for the date d considered, for the values of l corresponding to the days of the LT already elapsed, and for the values of w corresponding to prices equal to or lower than the price applied in these days.

This update can be based on: a) extrapolating the divergence, assuming it is attributable to systematic causes b) not extrapolating the divergence, assuming it is attributable to non-systematic causes c) assume that the divergence is so severe that it invalidates the existing forecast and a manual re-estimation is required. The user must set the parameters (B-1 in Figure 3) governing the percentage in which deviations would be extrapolated, as well as the threshold beyond which a manual re-estimation is required.

The last module, manual forecast update, offers the forecaster the possibility to adjust the forecast for specific dates as new information (other than the real booking data) becomes available, e.g. information on promotions, competitor's actions or fairs.

4. Conclusions

Given the very specific and exacting demand planning requirements of Yield Management based Decision Support Systems, and the present status of SME in that sector, the application of quantitative extrapolation based on the current historical data is likely to be impractical in many cases. The design proposed for the demand planning modules of such a DSS, developed in the framework of both a larger YM based DSS research project and a cooperative project with a small hotel chain, addresses the main stumbling block that have been identified, and provides a viable step towards the actual implementation of these DSS by the Spanish hotel SME. It proposes a gradual approach, whereby a first implementation can be based on a number of simplifying assumptions (e.g., assuming that the proportion of customers in each willingness-to-pay bracket is independent of the time remaining until the pernoctation date) that can then be progressively eliminated when the system can sustain the additional complexity involved. The proposed approach also opens two promising evolution paths: the combination of statistical extrapolation with human estimation (e.g. the forecaster might estimate the means of the factors, while as their correlations can be estimated from past data) and the dual applicability of the modelling and factorial decomposition involved both for forecasting and for simulation.

References

Baker, T.K.; Collier, D. (2003). "The benefits of optimizing prices to manage demand in hotel revenue management systems". Production and Operations Mgt., 12(4):502-518.

Boyd, E.A. (1998). "Airline alliance revenue management". OR/MS Today, 25:28–31.

Cavalieri, S.; Macchi, M.; Valckenaers, P. (2003). "Benchmarking the performance of manufacturing control systems: design principles for a web-based simulated testbed". Journal of Intelligent Manufacturing, 14(1):43–58.

Chiang, W.C.; Chen, J.C.H.; Xu, X. (2007). "An overview of research on revenue management: current issues and future research". Int. J. Revenue Management, 1(1):97-128.

Elliott, T.L. (2003). "Maximising revenue production while cutting costs: an airline industry mandate". Journal of Revenue & Pricing Management, 1:355–368.

Varian, H.R. (1996). "Differential pricing and efficiency". FirstMonday, 1(2):1-16. URL: http://www.firstmonday.org/issues/issue2/different/

Varian, H.R. (2007). Intermediate Micro Economics: A Modern Approach (7th ed.). W. W. Norton & Company.