Revenue Management Tutorial

Stefan Pölt

Network Management
Lufthansa German Airlines
stefan.poelt@dlh.de

AGIFORS Reservations & Yield Management Study Group
Berlin, 16-19 April 2002

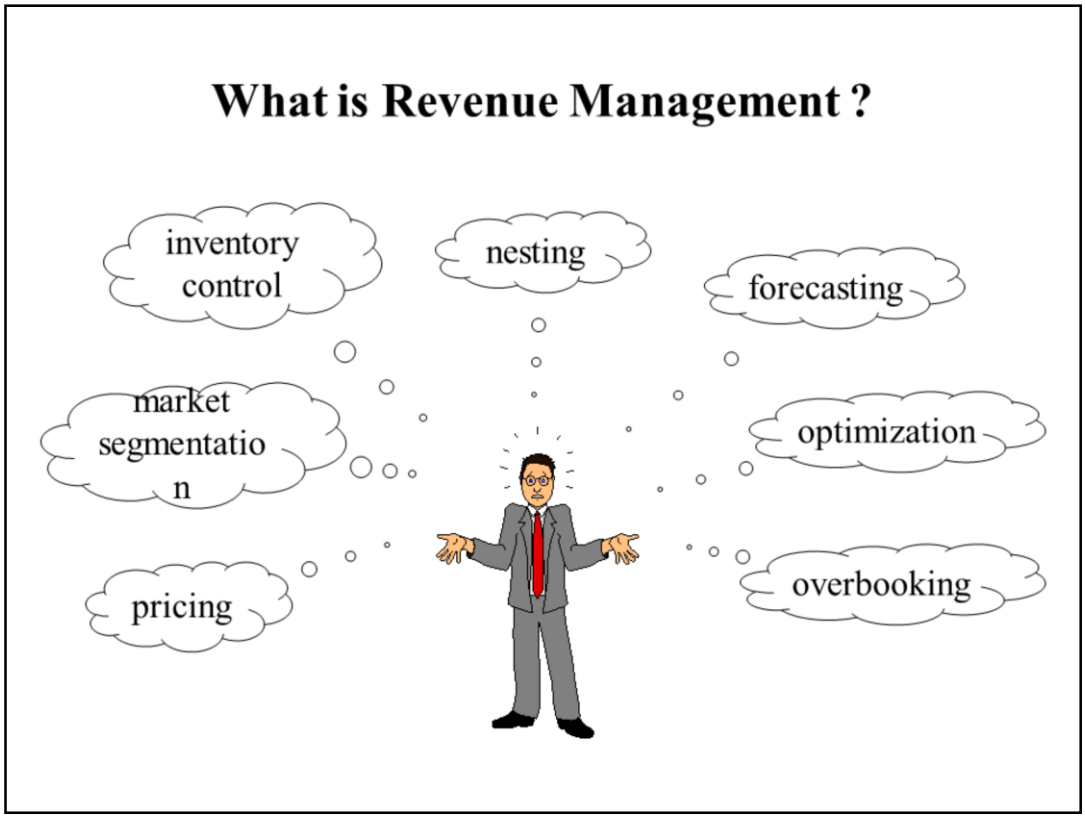

Let's try to clarify what Revenue Management means. If you ask 5 different people what's the scope of Revenue Management you will probably get 5 different answers.

What is Revenue Management?

CLICK

Is it inventory control?

CLICK

Is it forecasting?

CLICK

Is it optimization?

CLICK

What about market segmentation and Pricing?

CLICK

Some other terms that are related to Revenue Management are 'class nesting' and 'overbooking'. In the following I will try to explain what all these terms stand for.

Revenue Management Definitions

'Selling the right seats to the right customers at the right prices and the right time' (American Airlines 1987)

(Squeezing as many dollars as possible out of the customers)

'Integrated control and management of price and capacity (availability) in a way that maximizes company profitability

In the late 80s American Airlines gave the following popular definition of Revenue Management: 'Selling the right seats to the right customers at the right prices'. Later on this definition was extended by 'at the right time' and it will continue to be changed as Revenue Management is still under development.

So I have already seen the extension 'through the right distribution channels' to address the growing importance of internet bookings.

CLICK

The ultimate goal behind Revenue Management is to maximize the money that customers spent on us.

CLICK

Another, more general definition of RM is: The integrated control and management of price and capacity in a way that maximizes company profitability.

There are also other definitions like: Matching the demand to an existing supply by charging different prices and offering the optimal capacity for each demand category in such a way that the maximum revenue is generated.

Since Revenue Management is a complex field it is not easy to capture all aspects in a single definition.

Revenue Management History

- RM was 'invented' by major US carriers after airline deregulation in the late 1970's to compete with new low cost carriers
- Matching of low prices was not an alternative because of higher cost structure
- American Airline's 'super saver fares' (1975) have been first capacity controlled discounted fares
- RM allowed the carriers to protect their high-yield sector while simultaneously competing with new airlines in the low-yield sector
- From art to science: By now, there are sophisticated RM tools and no airline can survive without some form of RM
- Other industries followed - hotel, car rental, cruise lines etc.

Let's have a brief look at the history of Revenue Management .

Revenue Management was introduced by major US carriers as a reaction on new low-cost carriers started up in the late seventies after US airline deregulation.

The first reaction has been to match the low prices, but this was not successful because of the much higher cost structure of the big carriers.

One of the first Revenue Management instruments were the 'super saver fares' of American Airlines which have been the first capacity controlled discounted fares in the Airline market.

The principle of placing booking limits on discounted fares allowed the big carriers to protect their high-yield market segments while simultaneously competing with the new low-cost carriers in the low-yield segment.

In the meanwhile Revenue Management has become an industry standard with sophisticated tools in place.

The revenue gains from applying Revenue Management have been estimated between 10 and 30 per cent and no Airline will survive without some form of Revenue Management.

Other industries like Hotels, car rentals, cruise lines and so forth followed and adopted the Revenue Management principles to their needs.

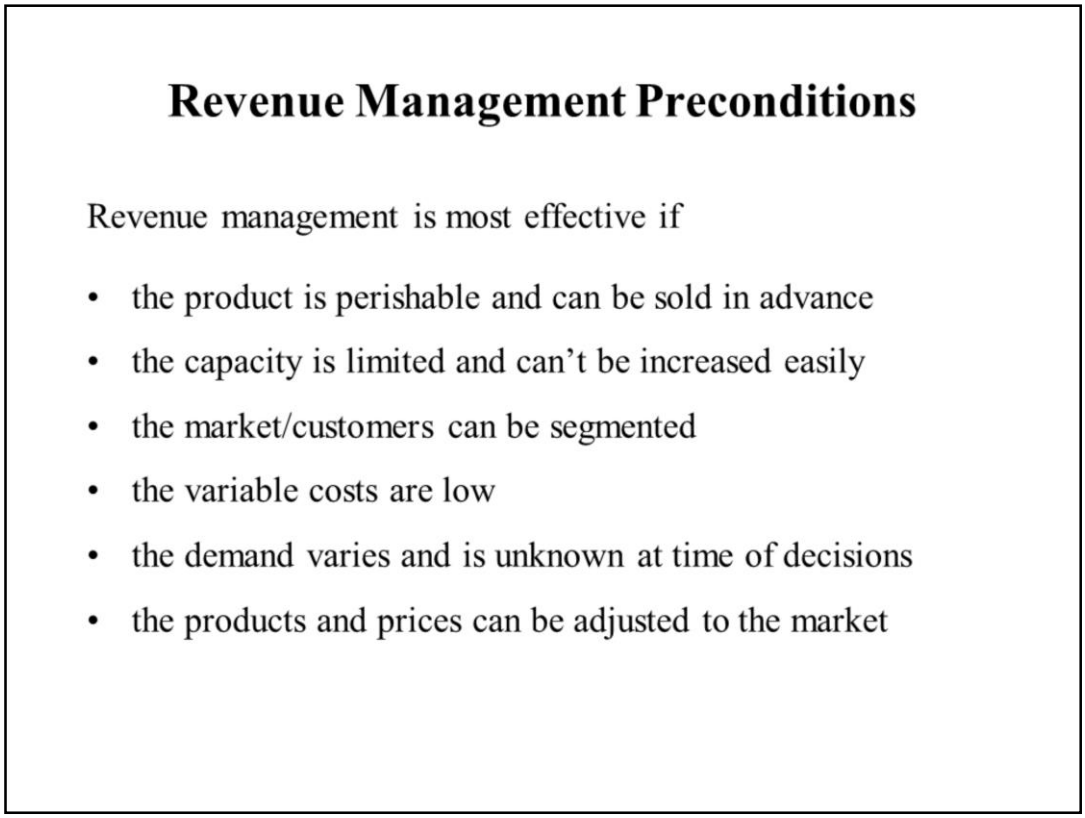

There are several characteristics of industries which make the application of Revenue Management most effective:

If the product is perishable, which means that it cannot be stored, at least not without significant cost or aging.

If the capacities are limited and and it's impossible or very expensive to increase capacity quickly.

If the market can be segmented. That means the customers are heterogeneous in terms of their needs and willingness to pay and can be split into price-sensitive and price-insensitive segments.

If the variable costs of a capacity unit are low. In the airline industry the variable costs consists of a reservation, a ticket, a bit more fuel, an additional meal and so forth, and are very low compared to the fixed costs of flying a plane from A to B.

If the demand varies. With a fixed and known demand the problem becomes much more easier and their is no need for sophisticated Revenue Management strategies.

If the products and prices can be adjusted to the market. One facet of Revenue Management is to react by pricing actions on changing market conditions. Markets change because of new competitors, changes in economy and so forth.

All of these characteristics hold for the airline industry. So it is no surprise that Airlines have been the inventors and first users of Revenue Management.

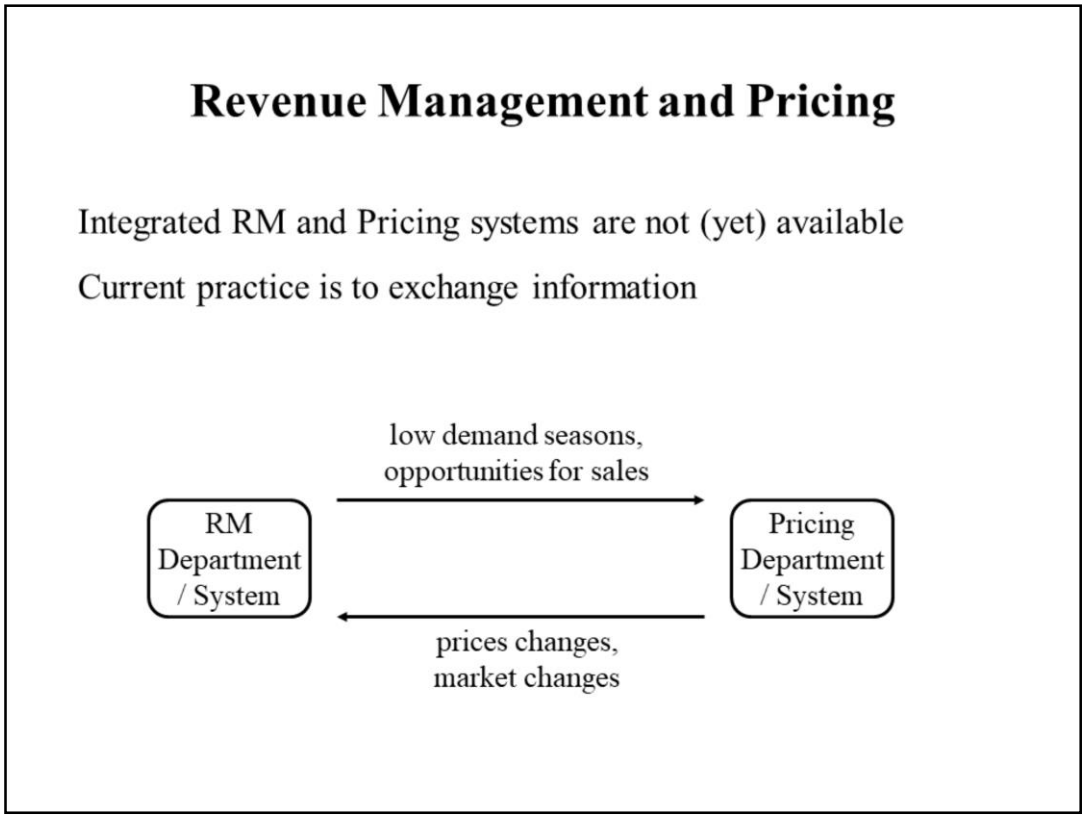

Revenue Management and Pricing are deeply interconnected. Sometimes Pricing is defined as a part of Revenue Management.

Unfortunately, on the system side there is no integrated Revenue Management and Pricing tool available. Usually Airlines have a sophisticated Revenue Management tool and a Pricing data base with automated price matching and distribution components.

The hard part in Pricing optimization is the estimation of price elasticity curves. It is not easy to gather the necessary data for the calibration of such models (in a good quality).

CLICK

Hence, at most Airlines there is only an exchange of data or information between the Revenue Management and the Pricing department. Revenue Management forecasts can be used in Pricing to indicate markets with low demand over a longer period where it might be a good idea to start a sale.

Pricing which is in general closer to sales and to the markets can tell Revenue Management general market changes and expected demand changes due to pricing actions. This information has to be fed back into the Revenue Management systems in order to adjust the forecasts of demand.

Revenue Management and Pricing

Goal is to adjust the demand to the 'fixed' capacity

Save seats for high-fare demand on full flights and channel low-fare demand to empty flights

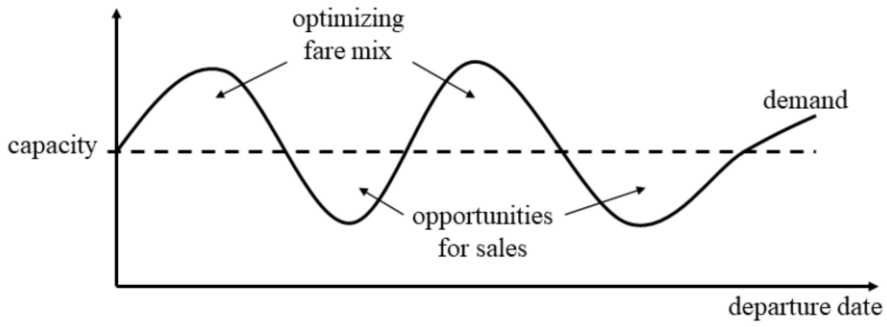

Since demand varies a lot over seasons, day of week and time of day, and since capacity is relatively fixed, it's a straight-forward idea to adjust the demand to the given capacities.

CLICK

For high demand flights there is an opportunity to pick the best bookings and spill the rest to the low demand flights that have enough capacity.

CLICK

For flights with low demand there is a chance to attract additional demand either by recapturing own spill, stealing demand from competitors, or by stimulating new demand by sales.

Market Segmentation

Passengers are very heterogeneous in terms of their needs and willingness to pay

A single product and price does not maximize revenue

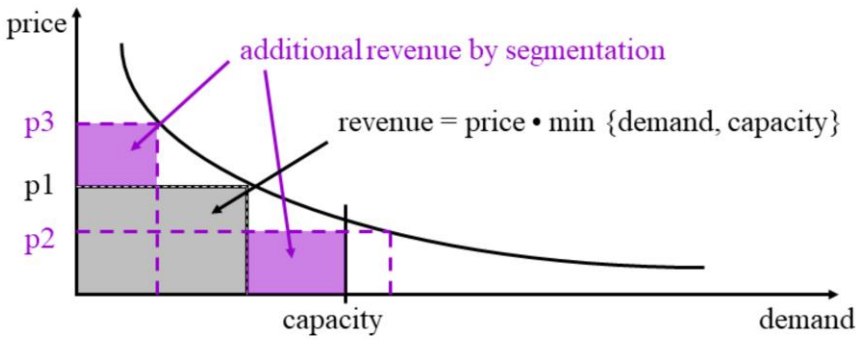

One of the preconditions for successfully applying RM to the airline industry is the heterogeneity of passengers in terms of their needs and willingness to pay.

Business travelers are usually time sensitive and are willing to pay a premium price for travel flexibility while leisure passengers are more price sensitive and are willing to change their travel plans in order to get a discounted fare.

This allows the airlines to define different products even within the same service level and to charge different prices for these products.

In the graph you can see a standard price-elasticity curve. The higher the price the smaller the number of people who are willing to pay that price. A single price ‘p1’ leads to a certain demand and the corresponding Airline revenue is the rectangle.

CLICK

With the introduction of two more prices you have the opportunity to generate additional revenue, the purple rectangles. But you have to prevent the people who are willing to pay the high price ‘p3’ from buying down to the lower prices ‘p1’ or even ‘p2’.

The most effective way to prevent buy down is conditioning. This means that low fare products are designed in such a way that not many business passengers are willing to accept their conditions.

Market Segmentation

- Products (booking classes) are different in
  - service (compartment)
  - conditions (advanced purchase, Saturday night stay, non-refundability etc.)
  - price
- Effective conditioning is essential for market segmentation (to prevent buy-down)
- RM is the last chance to mitigate the impacts of bad pricing decisions

The most effective conditions of low-fare products are

Saturday night stay - most business travelers want to return before the week-end.

Seven, fourteen, or even twenty one days advanced ticket purchase - most business travelers have to plan their trips on a short-time basis and have to be very flexible in changing their travel plans.

Tickets that are non refundable, endorsable, or re-routable are also less attractive to business passengers because they restrict their flexibility.

Some years ago Louis Busuttil told a funny anecdote about conditioning. At a Southwest Airlines flight from Dallas to Las Vegas the condition to get a special discount fare was to show-up dressed like Elvis Presley. So, it does not matter what conditions you put on discount fare as long as they effectively prevent high fare passengers from buying them.

Another aspect which prevents some business passengers from buying down is service (and maybe status symbol). Especially at long-haul flights the seat pitch and other comforts are an selling argument.

In the narrow sense, Revenue Management optimizes on given capacities and given prices. In this sense Revenue Management is the last chance to steer against bad pricing decisions by restricting the availability of low fare products.

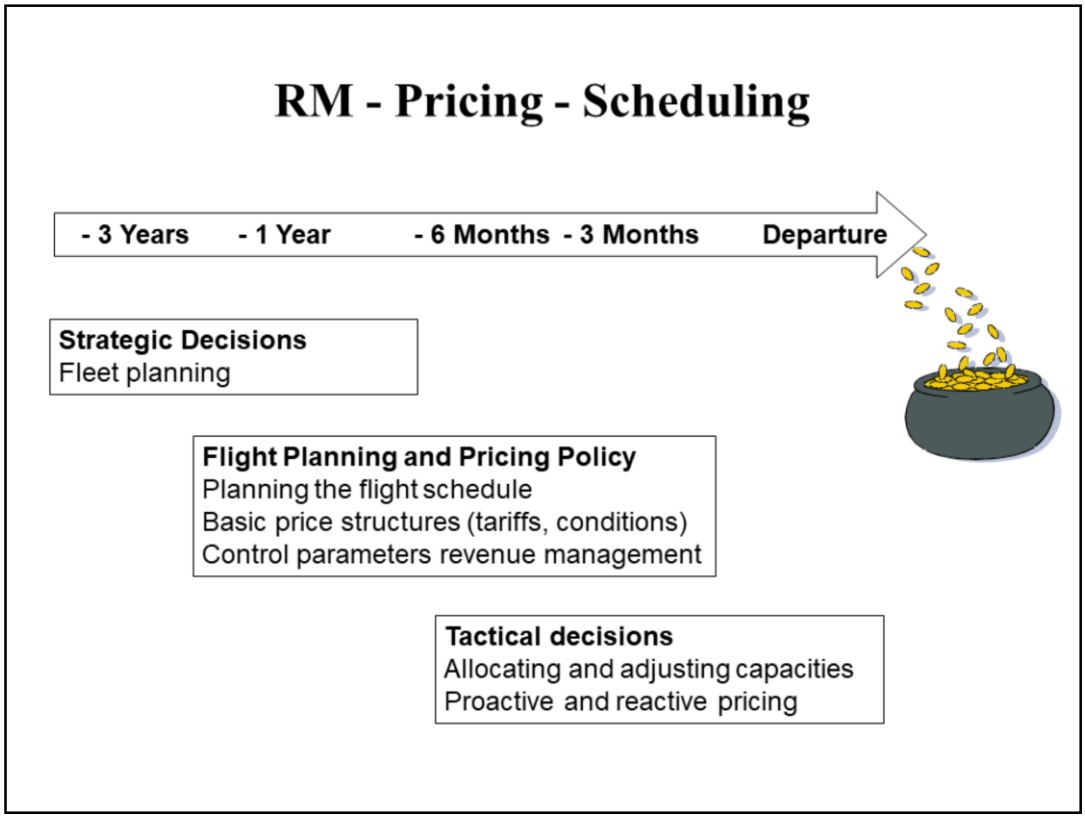

Although it is a too complex task to jointly optimize Revenue Management, Pricing and Scheduling decisions, there are deep interactions.

On a time before departure axis the first decisions within Network Management are strategic ones like Fleet Planning.

About one year prior to departure the flight schedule is planned and the core price structure and general control parameters within Revenue Management are determined.

Tactical pricing and inventory control, as well as short-time equipment changes usually happen within the last half year before departure.

The different time-lines and the different aggregation levels of data makes the integration of Scheduling, Pricing and Revenue Management difficult.

But, one could think at least of synchronized demand forecasts, starting in the long-term at a higher aggregation level and being refined as departure comes closer.

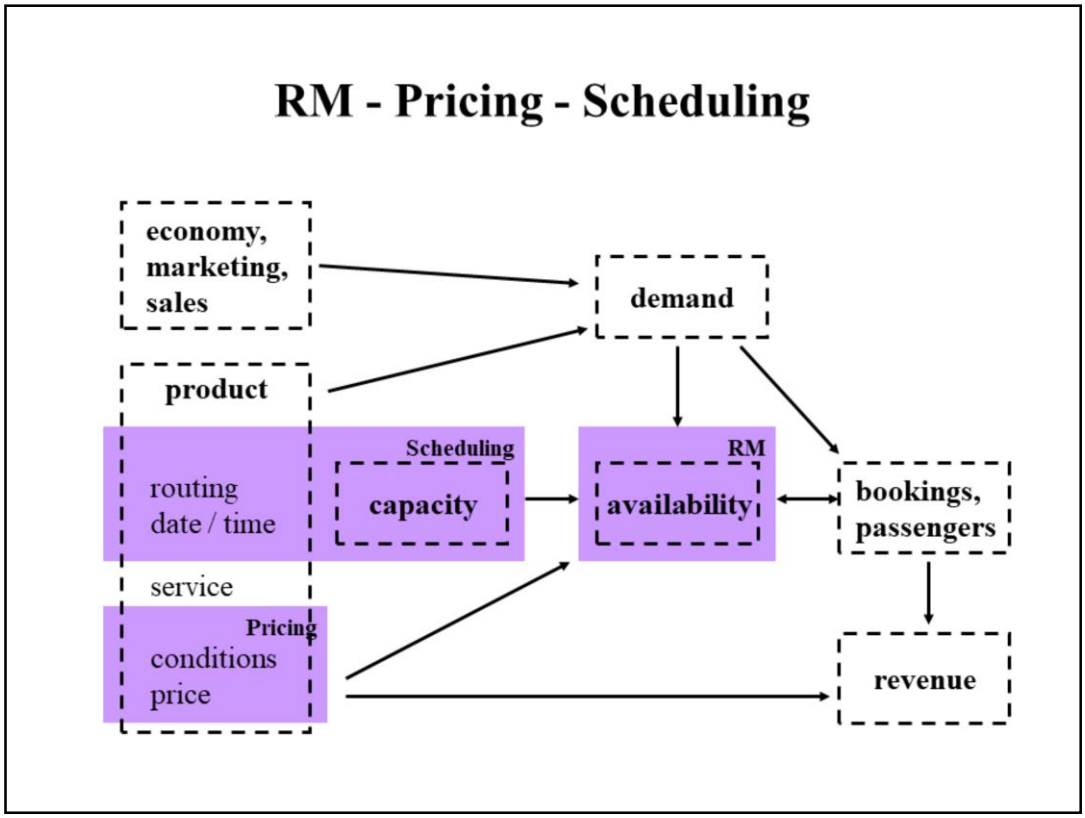

This slides gives another view on the interactions of Scheduling, Pricing and Revenue Management. All three together have a big influence on the total Airline revenue.

The product an airline offers is to a great extent defined by Scheduling and Pricing. Scheduling defines the routing, the frequency, the departure time, whether it is a non-stop or a connection. Pricing defines the price and the conditions. There are other features of the product like service, seat pitch, lounges and so on which are defined by product management and frequent flyer programs which are not mentioned here.

The quality of the product determines the demand for it. There are other external factors like economy, marketing and sales effort and so forth which also have an influence on the demand.

The role of Revenue Management is to match the demand with the capacities given by Scheduling. This is done by determining the availability of the product. In order to optimize the availability, Revenue Management has to know how much money the company will get when this product is sold. For this purpose either the fares from Pricing can be used - as shown here - or historical average revenues from Revenue Accounting.

The availability together with the demand define how much bookings or passengers you will have for each product. This contributes to the total revenue.

While this is a very simplified picture, it has a central message. All three together, Scheduling, Pricing, and Revenue Management have a great influence on the total Airline revenue. To maximize overall revenue the decisions within Scheduling, Pricing and Revenue Management should be harmonized.

Otherwise, you might have situations where Pricing starts a sale, but Revenue Management does not give enough availability to this product.

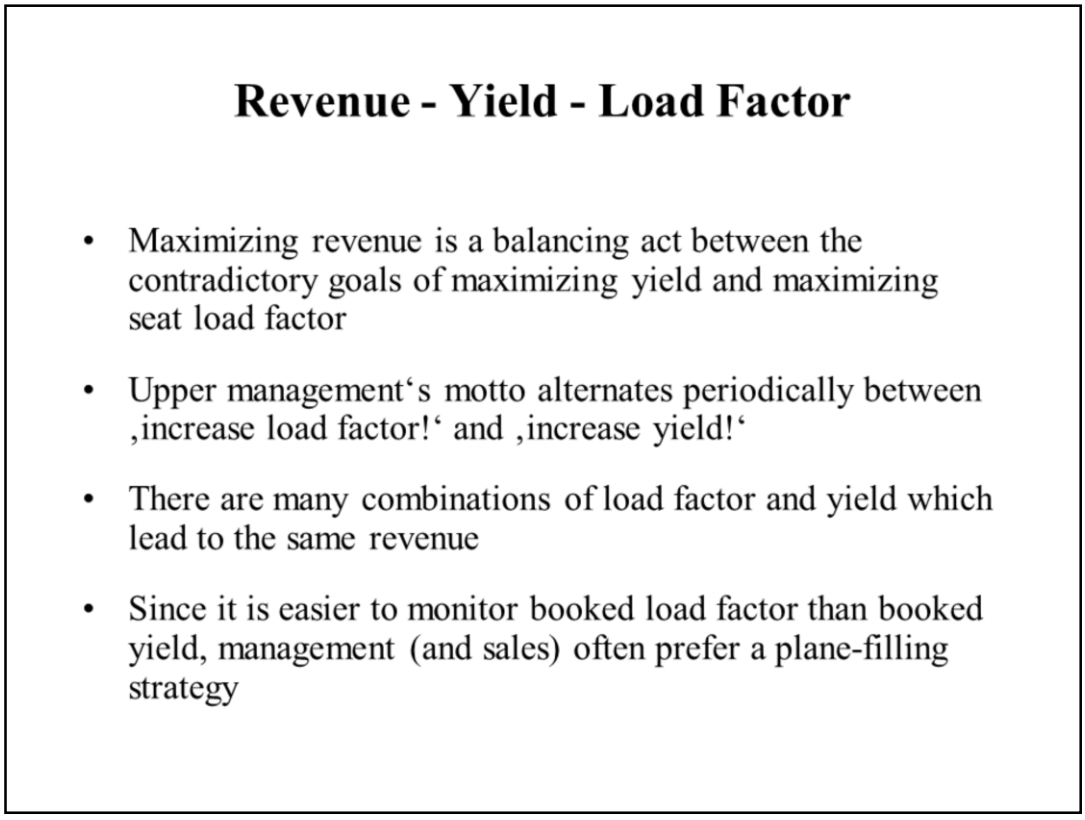

Maximizing the revenue is a balancing act between the two contradictory goals of maximizing yield and maximizing average seat load factor. Yield is defined as revenue per passenger or revenue per passenger mile.

While this AGIFORS study group has the name ‘Reservations and Yield Management’, the term ‘Yield’ is a bit misleading. We don’t want to maximize yield (only).

And we don’t want to maximize sales or seat load factor (only). An article of American Airlines notes:

If we only sell full fares we’ll be lonely and go broke.

If we sell everything at discount fares, we’ll be popular until we go broke.

It’s not an easy task to find the right middle-way. Since it is easier to monitor booked load factor the motto of the upper management is often ‘fill the planes’, especially in situations where the booked load factor is lower than last year. This motto lasts until an CEO gets a report that states a decrease of yield. Then the motto changes to ‘increase yield’. That is a periodical game similar to an electronic control element with a feedback loop that swings around an optimum state.

There is no single optimal combination of yield and seat load factor. There is a whole range of control strategies from restrictive to more open which might lead to the maximum or close to the maximum revenue. But, the more you reach the borders of this range the greater is the risk that your control gets sub-optimal.

Revenue Management Dilemma

- High-fare business passengers usually book later than low-fare leisure passengers
- Should I give a seat to the $300 passenger which wants to book now or should I wait for a potential $400 passenger?
- Most decisions in Revenue management are based on balancing risks, costs, or opportunities

The world of Revenue Management could be so easy if the full-fare passengers would book before the discount passengers (and if we wouldn't have any no-shows or cancellations).

Unfortunately, it is closer to the opposite. Most full-fare passengers are late booking business passengers and most discount passengers are early booking leisure passengers.

Here the game of Revenue Management starts. Should I give a seat to the passenger who stands in front of me and is willing to pay 300 dollars or should I gamble and wait for a potential 400 dollar passenger which might show-up later?

A similar risk-balancing game is played in overbooking. Should I increase the overbooking level by one, lowering the risk of empty, spoiled seats but simultaneously increasing the risk of over sales and denied boardings?

Overbooking

- Some (about 13% on average) booked passengers don't show-up at departure due to
  - double (fake) bookings
  - missed connections etc.
- Overbooking to compensate for no-shows was one of the first Revenue Management functionalities (1970's)

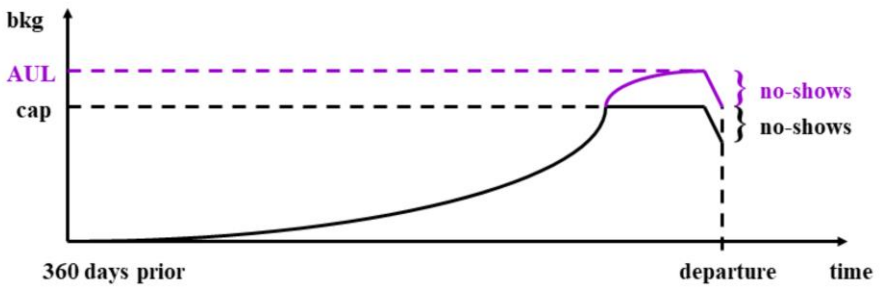

The reason for overbooking is that some booked passengers don't show up at departure - in average around 15%. This might be due to missed connections or double or fake bookings which have not been cancelled.

To allow more bookings than there are physical seats on the plane was one of the first Revenue Management functionalities back in the 70s.

The following graph shows for a high demand flight how the bookings would evolve over the booking period without overbooking. At some point of time the bookings reach the capacity and some booking requests are rejected. At departure there are no-shows and the plane departs with a substantial amount of empty seats.

CLICK

With overbooking more bookings than capacity are allowed.

CLICK

This example shows a perfect overbooking case which almost never happens due to all the uncertainty in no-show forecasts. In this case the overbooking exactly compensates for the no-shows and the plane departs with a hundred per cent seat load factor.

Overbooking

Sophisticated overbooking algorithms balance the expected costs of spoiled seats and over sales

Typical revenue gains of 1-2% from more effective overbooking

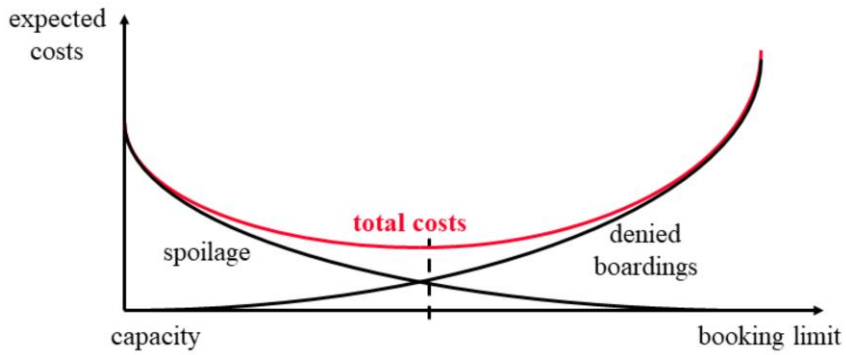

A simple overbooking algorithm takes the no-show forecast and overbooks to compensate for those no-shows.

A more sophisticated overbooking takes the different costs of no-shows and denied boardings into account as well as the uncertainty of the no-show forecasts. It calculates the expected costs of spoiled seats and denied boardings for each possible overbooking level and selects that one with minimum expected costs.

This graph shows the two cost types.

CLICK

The risk of spoilage, that is empty seats despite high demand, is the greater, the smaller the overbooking limit is.

CLICK

On the other hand the risk of denied boarding increases with increasing overbooking limit.

CLICK

The sum of both costs has a minimum and the corresponding booking limit minimizes the expected total costs. The literature reports revenue gains between one and two per cent by effective overbooking algorithms.

A problem with cost based overbooking algorithms is the estimation of spoilage and especially of denied boarding costs. It is not easy to estimate the loss of reputation associated with a denied boarding.

Hence the trend goes to service based overbooking algorithms where you can set an upper bound on the number of expected denied boardings or on the probability of

denied boarding events.

Upgrading

Fixed cabin capacities do no suit demand in all cases

Upgrading is a 'virtual' shift of capacity between cabins to allow more bookings in the lower cabin

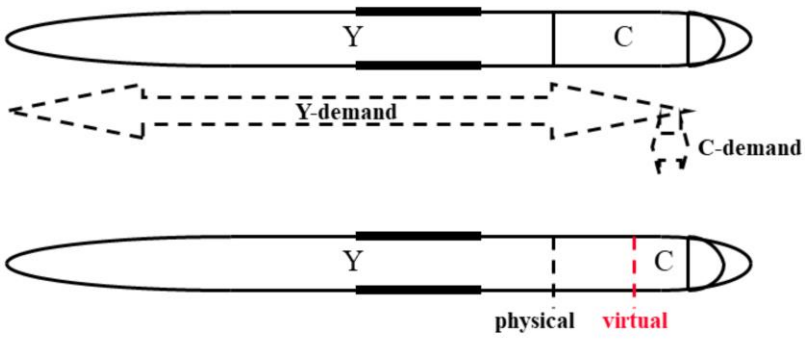

Another function of most Revenue Management systems is upgrading. The fixed cabin sizes of Intercontinental flights do not suit demands in all cases. You might have excess demand in Coach cabin and excess capacity in Business cabin. It is an airline policy question if you should virtually shift some seats of Business cabin to Coach cabin in order to satisfy more Coach demand.

CLICK

From a short-term revenue point of view you should, but such a policy has the risk of dilution. If upgrading happens on a regular basis, some passengers learn that there is a good chance of being upgraded for free and buy a Coach ticket instead of a Business ticket.

Forecasting

- Most important forecasting items in RM are
  - demand
  - no-shows
- Forecasting is usually based on historical bookings
- The mass of things to forecast makes automation (computer systems) necessary
- Systems allow influences to react on changing conditions that are not reflected in stored booking history (fare changes, competitors, special events etc.)

Forecasts are the basis for optimization in Revenue Management systems. The most important things to forecast are demand and no-shows or show-up rates. For management reports a forecast of passengers on board is interesting as well.

The forecasts are usually based on historical bookings and availabilities which are stored in a data base.

There is a mass of things to forecast. At Lufthansa, for example, we have about 1500 daily flights and 15 booking classes. This gives more than 40.000 forecasts. Since every flight is not forecast and optimized only once, the number easily explodes. Currently we produce several millions of forecast values per day.

This mass of data and forecasts cannot be handled without an automated Revenue Management system. Most of the flight events are not specifically considered by flight analysts. The motto is management by exception. The flight analysts concentrate on the most valuable peak flights (usually at holidays, fairs or other special events) where there is most money on the table.

Sophisticated Revenue Management systems allow the users to influence the forecasts at various aggregation levels in order to adjust them to changes that are not reflected in the booking history. There might be fare changes, changes in the market structure because an important competitor leaves the market, special events like Olympic games and many more.

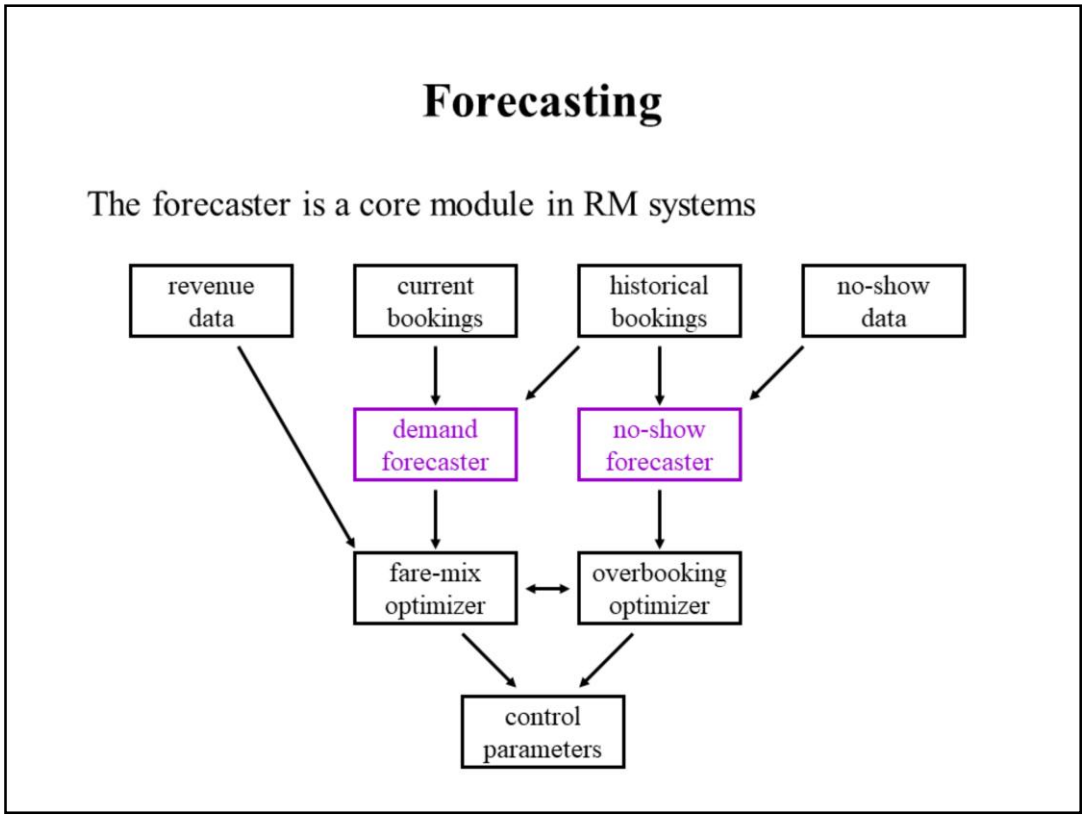

This picture shows the main modules of a typical leg or segment based Revenue Management system. The forecasts of demand are based on current bookings of the flight and on historical bookings of comparable flights. The no-show forecasts are based on historical bookings and no-show information which usually comes from a check-in system.

Both forecasts are used in the optimization. The no-show forecasts are used to calculate overbooking levels and the demand forecasts are used to calculate booking levels by booking or fare class.

The resulting control parameters are passed to the Computer Reservation system in order to control availability and booking requests.

Forecasting

- There are two possible consequences of bad demand forecasts: spoiled seats and bad fare mix (yield)
- As a rule of thumb, 10% improvement in forecast accuracy translates to 1-2% revenue increase
- If not covered by specific functionalities (sell-up, dynamic hedging, full fare/future protects) moderate over-forecasting increases revenue (especially at high-demand flights)
- There are two possible consequences of bad no-show forecasts: spoiled seats and denied boardings

There is a lot of variability and uncertainty in the demand forecasts, especially at the very detailed level at which Revenue Management forecasts have to be produced. Reasons are seasonality, fare changes, schedule changes, sell-up and diversion effects, spill and recapture, economical factors and so forth.

There are two possible consequences of bad demand forecasts: Empty or spoiled seats due to over-forecasting high fare demand and bad fare mix due to under-forecasting high fare demand.

As a rule of thumb, improving the forecast accuracy by 10 percentage points translates to a revenue increase of 1 per cent in average, on high demand flights up to 4 per cent.

It has been shown in several simulations that a moderate over-forecasting increases revenue especially on high demand flights, since it forces people to sell-up.

There are two possible consequences of bad show-up rate forecasts: Empty or spoiled seats due to over-forecasting show-up rates and over sales or denied boardings due to under-forecasting show-up rates.

Nesting

- Almost all Reservation systems allow serial/linear nesting of booking classes
- Nesting prevents high-fare booking classes being sold out when lower-fare booking classes are still open

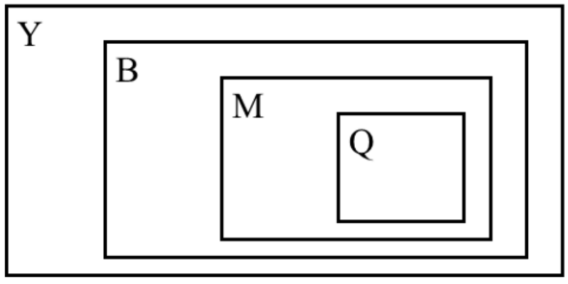

There are several ways to calculate and display availabilities in the Computer Reservation Systems. In the beginning, every booking class had it's own allocation and if the number of bookings had reached this limit, the booking class was sold out. This is known as partitioned control or discrete nesting.

The disadvantage of discrete nesting is that you might have situations where a high booking class is sold out while a lower booking class is still available.

Serial or linear nesting eliminates such illogical situations. In linear nesting there is a linear order of the booking classes and higher nested booking classes have access to the allocations of lower nested booking classes.

In this example, the allocation of highest Y-class contains the allocation of the lower nested B-class.

Simulations have shown that with realistic errors in the demand forecasts, linear nesting is the better method and results in revenue gains of 10 per cent and more.

Optimization

- Calculation of booking limits by booking class or Bid Prices
- EMSR robust and popular heuristic

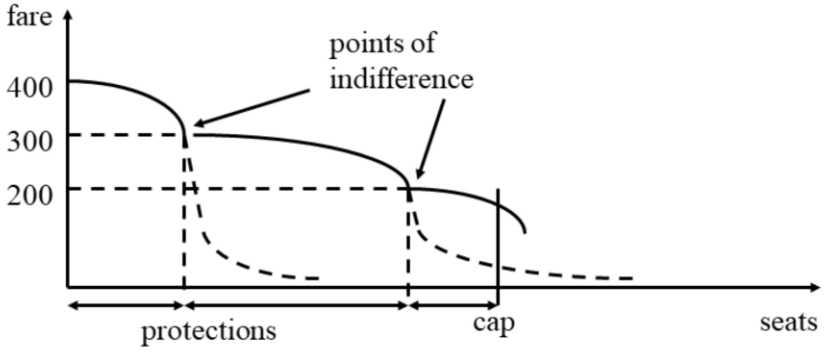

In fare-mix optimization the booking limits are calculated. A popular and robust heuristics for that step is EMSR (expected marginal seat revenue) published by Peter Belobaba in the late eighties.

It needs three different forecast values by booking class: mean demand, demand variability, and expected revenue or fare. EMSR gives an answer to the question posed at the Revenue Management dilemma slide: Should I accept a 300 dollar booking or should I gamble and wait for a potential 400 dollar booking?

In other words how many seats should be protected for the late booking 400 dollar passengers?

The forecasts of mean demand and demand variability can be used to forecast the probability that the demand is at least x passengers. This probability is decreasing with increasing x. With help of these probabilities one can calculate points of indifference where the expected revenue of protecting an additional seat to the 400 dollar booking class equals the expected revenue of giving this seat to the next lower 300 dollar booking class. This equilibrium is reached for that x where the probability of selling an additional 400 dollar ticket is down to 75 per cent which is the ratio of 300 and 400 dollars.

Leg - O&D

Leg control can't distinguish between local and connecting traffic

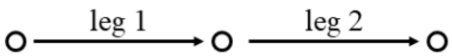

low demand	low demand	leg control sufficient
low demand	high demand	prefer connecting traffic by O&D
high demand	high demand	prefer local traffic by O&D control

O&D control can increase availability to long-haul passengers **AND** prevent long-haul passengers from displacing high-fare short-haul passengers

Leg-based Revenue Management methods like EMSR have a disadvantage – they cannot distinguish between one-leg and connecting demand. A connecting passenger gets availability for a booking class if and only if the class is open at both legs. And the optimization of booking class limits is done at both legs independently and separately.

Because of the fare structure of Airlines two one-leg passengers, one travelling leg1 and the other travelling leg2, usually give more revenue than one connecting passenger, while in both cases the same number of seats on the same flights are used.

If there is low demand on both legs, the independent optimization of legs does not hurt since every booking request - regardless of one-leg or connecting - should be accepted.

If there is low demand at one leg but high demand at the other, connecting passengers should be preferred to one-leg passengers on the bottle-neck leg. But, a leg-based Revenue Management system cannot distinguish between them.

If there is high demand at both legs, two one-leg passengers should be preferred to one connecting passenger of the same booking class. Again, leg-optimization cannot distinguish and accepts within a booking class in a first come first serve order.

Another advantage of O&D control is that it can distinguish different point of sales. For some countries it makes a big difference for the Airline at which place the booking was made and in which currency the ticket is paid.

O&D

- Slogan in RM during the last years
- Preconditions
  - structural (connecting traffic)
  - technical (seamless link to CRS, O&D data base, etc.)
  - organizational (market oriented RM organization)
  - soft factors (management commitment, intense training)
- Obstacles
  - complexity (complex algorithms, mass of data, source of errors)
  - data quality (dirty PNR data etc.)
  - costs (seamless availability, hardware, etc.)
  - cheating (need of 'married segments' and 'journey data')
- But, around 2% increase of revenue!

O&D control was THE slogan in Revenue Management over the last years and many major Airlines have already taken the step from leg to O&D or are considering it in the near future.

There are some preconditions to O&D control. The network should have substantial amount of connecting traffic. Otherwise the advantages of the refined O&D control don't take effect very often.

Since the GDS (global distribution systems) like SABRE, AMADEUS, GALILEO and so forth are leg or segment based, you have to have a seamless link to them in order to get the whole itinerary information which has to be evaluated in O&D control.

The Revenue Management department has to be re-organized from a leg-oriented to an O&D- or market-oriented structure.

And you have to have the commitment of the upper management to get the money and the power for all the necessary invest and re-structuring.

Some arguments against O&D control are increasing complexity, decreasing data quality at this very disaggregated level, increasing cost for building up and maintaining the seamless link and huge O&D data bases, and 'holes' in the system. You should prevent sales agents from playing around whether they can get better availability by booking two separate legs instead of a connect or by booking a connect and cancel one leg afterwards. Otherwise you can realize only a part of the O&D gains.

But after the dust has settled, there is a big argument for this step. It has shown in several simulations that O&D control increases revenues two per cent and more at realistic average demand and seat load factors. And it allows a better integration of RM with other parts of the airline, as Pricing and Scheduling.

... And Many More

- Group booking control
- Point of sale control
- Revision of RM decisions
- Data quality, outlier handling
- Reporting, monitoring, performance measurement
- Mathematics, algorithms, models etc.
- ...

There are a lot of other aspects within Revenue Management which I have not addressed in this presentation. For example, control of group bookings, point of sale control, the frequency of re-forecasting and re-optimization, the impacts of data and forecast quality, outlier detection, and the importance of monitoring and performance measurement.

Thanks for your attention. Are there any questions?