



Advances in revenue management: the last frontier

B. Vinod¹

Received: 20 April 2020 / Accepted: 20 August 2020 / Published online: 10 September 2020
© Springer Nature Limited 2020

Abstract

After deregulation of the airline industry, revenue management has evolved from flight leg and flight segment-based inventory controls to origin and destination inventory controls. Advances in revenue management have always focused on more granular controls to maximize revenues. Revenue management of individual seats is the most granular level of inventory control by customer segment or individual customer. This paper discusses the evolution of revenue management leading up to the inventory control of individual seats.

Keywords Seats · Seat types · Revenue management · Seat maps · Seat map cache

Historical perspective

The origins of yield management (or “revenue management” as it is called today) can be traced to Littlewood (1972) from the British Overseas Airways Corporation (BOAC) who proposed that airlines maximize revenues instead of passenger occupancy on a flight for the perishable seat inventory. Known as Littlewood’s rule for two booking classes, this can be extended to multiple classes. The dismantling of government control with deregulation of the US airline industry in 1978 resulted in an exponential increase in fare filings and new markets that were served by airlines. In the fiercely competitive environment after deregulation, several airlines like Braniff Airlines (1982), PEOPLExpress (1986), Eastern Airlines (1991), and Midway Airlines (1991) ceased to operate. There was also a wave of airline consolidations after deregulation.

The first yield management system was created in 1985 by American Airlines under the leadership of CEO Robert L Crandall to counter the competitive threat from PEOPLExpress. An Operations Research team under Tom Cook, President of American Airlines Decision Technologies, developed the system which was deployed in 1986 (Crandall 1995; Smith et al. 1992). The inventory controls for this system was based on leg/segment controls—leg class

nested inventory controls with segment close indicators and segment limit sales by booking class. Yield management was renamed revenue management in the 1990s since revenue and not yield is maximized. Revenue management is based on the fundamental premise that for the management of perishable inventory all customers are not created equal. As a result, airline customers are conditioned to paying different amounts for the same product depending on when they book and the details of the itinerary.

Leg/segment revenue management was followed by *Origin and Destination* inventory controls for network carriers to control the flow of connecting traffic through their hub airports. The first O&D yield management system based on virtual nesting controls was deployed in 1987 for American Airlines (Smith 1986; Smith et al. 1992; Vinod 1989). While a few airlines deployed virtual nesting in the 1990s such as United Airlines, Delta Air Lines, KLM Royal Dutch Airlines, and Scandinavian Airlines System (SAS), most network carriers began adopting continuous nesting controls (also known as bid price controls) (Vinod 1995). American Airlines and US Airways migrated to Sabre’s bid-price-based O&D system in 1998. The bid price is the opportunity cost of not having an incremental seat on a flight. Adoption of bid price controls by network carriers started in the late 1990s and continues to this day.

The financial impact of yield management to American Airlines was estimated by Crandall in 1992 to be \$1.4 billion in incremental revenue in the last 3 years (Smith et al. 1992). By 1998, Tom Cook had increased the estimated impact to “almost \$1 billion in annual incremental revenue” (Cook

✉ B. Vinod
benvinod@yahoo.com

¹ Sabre Research, 3150 Sabre Drive, Southlake, TX 76092, USA



1998). Revenue management is one of several applications that was classified as strategic operations research (Bell et al. 2003) since it creates a sustainable competitive advantage.

The growth of Low-Cost Carriers (LCCs) in the mid-to-late 1990s led to the creation of restriction-free tariffs (Gorin and Belobaba 2004; Vinod 2005; Fiig et al. 2005). Creation of restriction-free tariffs causes a problem with one of the fundamental assumptions of revenue management that demand for a booking class is independent of the other classes. With restriction-free pricing, price is the only determinant of the market segment and the demand for a booking class is therefore dependent or contingent on the lower class being closed. This led to the creation of new revenue management models to address restriction-free pricing, initially based on business rules and later with more sophisticated price demand curves. Network carriers who competed with LCCs in domestic markets with restriction-free tariffs and with other network carriers on international routes with regular restricted tariffs made extensive modifications to their revenue management models to operate in a hybrid mode due to the co-existence of regular tariffs and restriction-free tariffs in the airline's network.

To overcome the fundamental shortcoming of traditional revenue management, which determines inventory controls based on historical data and without considering competitor selling fares (fares and availability) in a market, *Competitive Revenue Management* makes tactical adjustments to inventory controls based on real-time shopping data (Ratliff and Vinod 2005; Vinod 2016). To achieve the strategic objectives by market, monitoring availability of selling fares in target markets is used to determine the probability that an itinerary will sell in order to determine the overrides to the inventory controls based on prevailing market conditions, to raise or lower availability by booking class. Leg/segment-based competitive revenue management was followed by extensions to O&D controls for network carriers, frequently referred to as *Dynamic Availability*. The earlier models were business rules driven followed with session-level optimization models that either maximized expected revenue or maximized net contribution. The attractiveness of an itinerary can be determined with a consumer choice model that is calibrated from a shopping request and response dataset based on pertinent variables such as displacement time, elapsed time, fare, screen presence, etc. Since *Dynamic Availability* is an inventory control recommendation to promote upsell in non-competitive situations, it requires no changes to existing systems and business processes.

Dynamic Pricing is closely related to dynamic availability. Both techniques can leverage competitive selling fares to arrive at an inventory control or dynamic price recommendation. The session-based fare optimizer already determines the optimal price point for the host airline based on the competitive set and current selling

fares of competing airlines in the marketplace. Instead of converting the optimal price point to an inventory control recommendation, the dynamic price is used to approximate the ticketed price. There are two versions of dynamic pricing—laddered pricing with discrete price points between published tariffs that are essentially private tariffs and continuous pricing. While dynamic pricing of opaque travel products (Zouaoui and Rao 2009) has minimal business process changes, the potential impacts of deploying dynamic pricing on third-party systems (e.g., GDS) can be quite significant (Choubert et al. 2015).

Prior to dynamic pricing, an itinerary is priced based on booking class availability and the applicable fares for the itinerary. During the itinerary pricing process, the total bid price for the itinerary is not considered as a hurdle price or minimum acceptable threshold and hence the ticketed price for the itinerary can be below the total bid price for the itinerary. Dynamic pricing overcomes this limitation by bridging the chasm between airline inventory and airline pricing by ensuring that the ticketed fare is greater than the total bid price for the itinerary. An added benefit of dynamic pricing is that it provides an infinite number of price points on the price demand curve that can be inventory controlled to generate incremental revenues.

Global airline ancillary revenues in 2019 exceeded 109.5 billion dollars (Ideaworks and Cartrawler 2019), which is a fivefold increase in ancillary revenues reported in 2010 of \$22.6 billion. *Offer Management* is the process of selling the right bundle of base airfare and air ancillaries to the right customer at the right price at the right time. Offer management extends traditional revenue management process of optimizing allocations for the base fare to include air ancillaries offered by an airline. At its core, offer management begins with customer segmentation based on context for travel using unsupervised learning techniques, followed by a recommendation engine that recommends offers for each customer segment and an offer engine to personalize and price the offer for a segment of one (Vinod 2017; Vinod et al. 2018). Model accuracy of a machine learning-based recommendation model, like a statistical model, does best until it is deployed since the statistical properties of the recommendation engine that was trained on an existing dataset with known variable ranges change over time due to unforeseen circumstances. To address the phenomena of *concept drift*, a test and learn experimentation model using reinforcement learning, frequently referred to as a multi-armed bandit, can continuously adapt the product recommendations to changing consumer behaviors, new entrants, and new product offerings from existing competitors. Following the lead of non-travel related companies like Netflix and Amazon, mass customization is touted as the way of the future in travel to create a unique offer based on the context for travel for a specific named customer. Acceptance of personalized



offers is an area of debate and remains to be seen if it will gain traction in travel.

The last frontier

Revenue management of seats—by seat type for a section of the aircraft or individual seat is in the not-so-distant future. When this happens, current static pricing (either based on rules practiced by the LCCs or based on fare components of the itinerary practiced by network carriers) of seats will be replaced by dynamic pricing of seats.

Determining availability and price by individual seat when the request is made is considered as the holy grail of revenue management. To revenue manage an itinerary at the seat level requires a real-time detailed seat inventory control component together with a revenue management capability to determine how individual seats or seat types should be priced.

In the current environment, based on the section of the airplane, seat prices are static by product (e.g., premium economy, economy, main cabin). The opportunity for variable seat pricing depends on how customers value a specific seat on an airplane. The dynamic price for a seat is influenced by the seat type (aisle, window, exit row, premium economy (wider pitch), basic economy), length of haul (short, medium, long), market, schedule attributes (aircraft type, departure time of day, non-stop versus connection), number in party (seats required together), channel, etc. These factors influence the probability of purchase of a specific seat. Availability by seat type also has an impact on air shopping. Online Travel Agencies (OTAs) display hundreds of flight options in a single shopping request. They also desire to display seat map counts by seat type based on user preferences to display with the itineraries returned as part of the shopping response. To avoid the computational burden, an approach is to deploy a seat map cache.

For airlines hosted on a specific Passenger Service System (PSS), the seat map cache is a combined PULL and partial PUSH model. PULL implies that a seat map service is requested to populate the seat map cache. A PUSH implies that every time the seat map is touched by a travel agent or end user, an update is sent to the seat map cache. In other words, PUSH is organic while PULL is inorganic. For example, if a carrier is hosted on a specific PSS such as Sabre, then the partial PUSH takes place only when a Sabre travel agent requests a seat map for a non-hosted carrier. When GDS subscribers of Travelport and Amadeus request a seat map, there is no mechanism to update the seat map cache in the Sabre PSS without active collaboration between all parties including the airline.

Consider the through flight A–C with legs A–B and B–C (Fig. 1).




Fig. 1 Through flight A–C

There are two *necessary and sufficient conditions* that should always be satisfied to ensure consistency and accuracy of seat map counts.

Condition 1 If a seat is consumed on a through flight (e.g., A–C), then the same seat should be consumed on the legs that make up the through flight (A–B and B–C).

Condition 2 If a seat is consumed on a leg, it should not be available on a through flight that includes this leg.

The seat map cache should maintain counts of sold and available seats along the following dimensions in real time: total seats, aisle/window/center seats, exit row seats, premium seats, preferred seats, no-charge seats, bulkhead seats, pay-for seats, seats together for a specific party size, and details for a specific seat (e.g., seat 12 A—window, premium, etc.). A cache controller is also required to determine when an item in cache needs to be refreshed. The refresh frequency can be based on static pre-departure reading day concept from revenue management or dynamic based on actual activity. In addition, the cache should also be able to display the physical seat map in lieu of asking queries.

Deployment of a seat map cache should support all the Passenger and Airport Data Interchange Standards (PADIS) seat characteristics an airline chooses to send. There are over a 100 characteristics, which can be grouped into the categories such as seat location (e.g., front of cabin, upper deck, adjacent closet, etc.), missing seats (e.g., no seats because of exit door, no seat because of upper stairs, etc.) seat characteristics (window, aisle, etc.), seat occupation details (e.g., occupied, advanced boarding pass issued, etc.), and seat blocking details (e.g., blocked for airport, blocked for through passenger, etc.).

Another important consideration during the sales process is seat-led shopping. Here are a few scenarios for seat-led shopping.

Show me flights to San Francisco, departing on June 4 before Noon and returning June 8 in the evening for a family of 4 and we want to sit together.

I want to go to London for 1 week, departing June 1. Only show me flights where an aisle seat in First Class is available.

I want to go to NYC on June 1, returning on June 5. Show me flights where exit row seats are available. In addition, I want to avoid regional jets.



Seat selection today is performed after the itineraries for display—schedule and price, are determined and the itinerary has been selected. This is a shortcoming since the seat map may display available seats that are not acceptable to a traveler forcing the traveler to start-over with a new flight search. Post-processing itineraries returned from shopping to select itineraries that fulfill the customer request are also not desirable since none of the itineraries returned may fulfill the seat request constraint. Seat-led shopping as a post-process of the shopping process (Sabre 2016) has its limitations as shown in the prototype demonstration. Ideally, seat-led shopping should be in-path in the shopping algorithm and not a post-process. This implies that the seat selection constraint imposed by the customer is considered *apriori* when schedules are generated by the shopping algorithm. For example, if a customer wants four seats together, the only itineraries that should be displayed are those which have four seats together that can be selected, although the prices for the itineraries may be higher without the seat type constraint. Seat-led shopping advances the user experience to select itineraries that guarantee the requested seat type request at time of shop. The seat map cache is used to generate schedules in the shopping algorithm to fulfill the seat requirement.

If seat pricing is by leg, there is an opportunity to incorporate a variable pricing concept like a bid price curve that makes the seat price as a function of the seats sold for a specific type of seat. This will require the airline's inventory control system to maintain seat counts by seat type by flight leg and date.

Consider the example ancillary pricing structure for *aisle* seats, for booking class M for a specific flight leg by departure date shown in Table 1

When availability is returned by an airline's inventory system, booking class combinability (ATPCo Category 10) that is determined during shopping is not yet known. Hence, once the itineraries are priced by shopping based on the reservations booking designator (RBD) availability, a second pass is required wherein the itineraries are sent to the airline inventory system for a total seat price by itinerary. This information will be used to determine the top N itineraries in the shopping response. Assuming the pricing of seats is by flight leg (segments and market seat prices can be derived from the flight leg seat prices), this information can be stored on the flight detail record. An issue

Table 1 Price as a function of seats sold

Number of seats sold (aisle seats)	Price
1–5	\$15
6–10	\$18
10–15	\$22
> 15	\$25

that needs to be addressed is determining the *total itinerary price with seats* since it is possible that a more expensive RBD may have seats at no cost and shopping would miss this RBD in the first pass. A push-back mechanism may be required to address this issue. Determining seat prices can be rule-based or based on advanced decision support that generates these price points for ancillaries by flight leg and date. Sophisticated methods can also be used to determine seat pricing. Machine learning methods such as logistic regression, gradient boosting model, and neural networks can be calibrated to estimate the accuracy, sensitivity (recall or true positive rate), precision, and false-positive rates. Monetizing seats with a dynamic seat price are an active area of research by leading airlines today.

With this approach, ancillary seat prices are dependent on total sales by seat type. In this scenario, an ATPCo OC filing for an aisle seat will only have a *reference fare* by market and time of day (if the airline participates in ATPCo) for informational purposes, and then the ancillary price for a seat will be a function of the bookings by seat type and potentially type of customer—the higher the seats sold count by seat type in this scenario, the higher the price. In this scenario, all sales channels *will* have to go to the airline (host CRS) to price the ancillary services. This is also a requirement for IATA's New Distribution Capability (NDC).

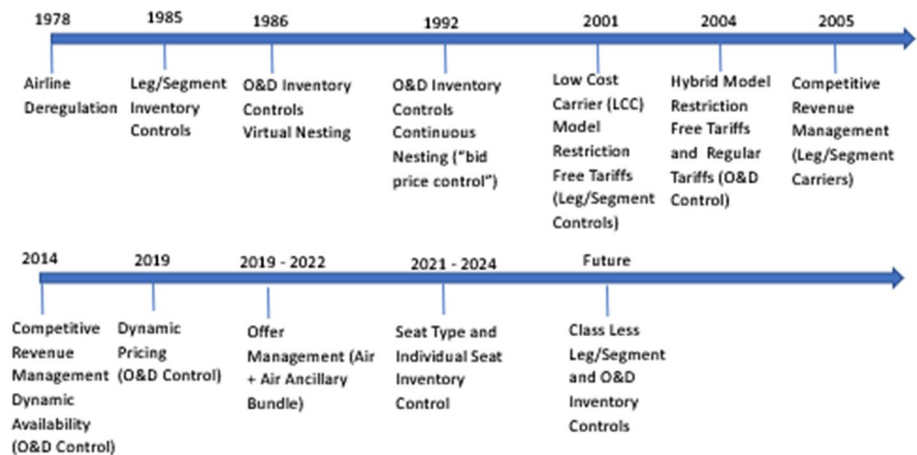
Impact of NDC on revenue management

In today's environment, a travel agent subscribing to a GDS can shop, book, price, and ticket an itinerary for a customer. With IATA's New Distribution Capability (NDC) that is currently being rolled out, pricing power shifts from the GDS to the airline. Hence, an agency must request itineraries and prices from an airline in an NDC world and hence implementation of this model and its variations will be within the domain of an airline's environment. The promises of NDC and the new XML-based messaging standard between airlines and GDSs are several. They are as follows:

1. Enable an airline to differentiate and offer products and services to customers that differentiate from a competitor.
2. Enable airlines to offer personalized offers with rich content that is not available in the GDS. Personalized offers can increase revenues at time of booking with ancillary sales such as pre-reserved seats, baggage, wi-fi, lounge access, meals, etc.
3. Maximize revenues with dynamic pricing and dynamic ancillary bundles based on context for travel.
4. Reduce the cost of distribution.



Fig. 2 Evolution of revenue management



Adoption of NDC is not without its challenges. First, when pricing power shifts from the GDS to the airline, a far greater investment is required by airlines in compute power and software to respond to every request from a travel agent. Second, scalability is a concern and it remains to be seen if NDC can scale to current GDS volumes for transaction processing without resorting to a cache-based solution (Vinod and Huff 2019). Third, is the processing of interline booking requests.

Another bi-product of NDC in the future is class-less revenue management. Revenue management without information on availability by booking class is a radical departure from the current environment (Isler and D'Souza 2009; Isler 2016). Booking classes are used today to distribute availability status using availability status (AVS) and direct connect between an airline and a GDS. In an NDC world, since an airline processes all requests, booking classes are technically not required since the GDS does not book and ticket a customer's booking request.

Figure 2 summarizes the evolution of revenue management in the airline industry since deregulation.

Conclusions

Revenue management has evolved along three dimensions. The primary dimension is greater level of granularity in inventory controls from leg/segment controls by booking class to O&D controls by booking class to control of inventory at a seat level. The second dimension has been focused on competitive revenue management with Dynamic Availability and Dynamic Pricing capabilities. The third dimension is offer management to maximize revenues by controlling the availability and price of a bundle consisting of the base fare and air ancillaries. Given its strategic importance, travel suppliers that manage perishable inventory will continue to invest in this application in the years ahead.

References

- Bell, P.C., C.K. Anderson, and S.P. Kaiser. 2003. Strategic operations research and the Edelman Prize finalist applications 1989–1998 [Electronic version]. *Operations Research* 51 (1): 17–31.
- Choubert, L., T. Fiig, and V. Viale. 2015. Amadeus dynamic pricing. In *AGIFORS revenue management Conference presentation*, 15 May, Shanghai, China.
- Cook, T. M. 1998. SABRE Soars. *OR/MS Today* (June) 26–31.
- Crandall, R.L. 1995. The unique U.S. Airline Industry. In *Aviation Daily's handbook of airline economics*, ed. D. Jenkins. The Aviation Weekly Group of the McGraw-Hill Companies, September.
- Fiig, T., K. Isler, C. Hopperstad, and R. Cleaz-Savoyen. 2005. *Davnmr: A unified theory of O&D optimization in a mixed network with restricted and unrestricted fare products*. Cape Town, South Africa: AGIFORS Revenue Management and Distribution Study Group Meeting.
- Gorin, T., and P. Belobaba. 2004. Revenue management performance in a low fare airline environment: Insights from the passenger origin-destination simulator. *Journal of Revenue and Pricing Management* 3 (3): 215–236.
- Ideaworks Company.com and Cartrawler. 2019. Cartrawler worldwide estimate of Ancillary Revenue for 2019. <https://www.cartrawler.com/ct/ancillary-revenue/worldwide-ancillary-revenue-2019>.
- Isler, L., and E. D'Souza. 2009. GDS capabilities, OD control and dynamic pricing. *Journal of Revenue and Pricing Management* 8: 255–266.
- Isler, K. 2016. *Revenue management in a world without booking class availability*. ATPCO Workshop, Washington DC, April 23.
- Littlewood. 1972. Forecasting and control of passenger bookings. Proceedings of 12th AGIFORS symposium proceedings, Nathanya, Israel. Reprinted in *Journal of Revenue and Pricing Management*, Vol. 4, 2005.
- Ratliff, R.M., and B. Vinod. 2005. Airline pricing and revenue management: a future outlook. *Journal of Revenue and Pricing Management* 4 (3): 302–307.
- Sabre. 2016. Seat led shopping. <https://www.sabre.com/insights/innovation-hub/prototypes/seat-led-shopping/>
- Smith, B.C. 1986. O&D control with virtual nesting, Internal Technical Report, American Airlines.
- Smith, B.C., J.F. Leimkuhler, and R.M. Darrow. 1992. Yield management at American Airlines. *Interfaces* 22 (1): 8–31.
- Vinod, B. 1989. A partitioning algorithm for virtual nesting indexing using dynamic programming. In *Internal Technical Report*, Sabre Technology Solutions, March 1989.



- Vinod, B. 1995. Origin and destination yield management. In *Aviation Daily's handbook of airline economics*, ed. D. Jenkins. New York: McGraw Hill.
- Vinod, B. 2005. Fare'ly simple. *Ascend* 19–21.
- Vinod, B. 2016. The evolution of yield management in travel. *Journal of Revenue and Pricing Management* 15 (3–4): 203–211.
- Vinod, B. 2017. The evolving paradigm of interactive selling based on consumer preferences. In *21st century airlines: Connecting the dots*, 207–213, September. ed. Nawal Taneja.
- Vinod, B., C. Huff. 2019. How to scale NDC to GDS transaction volumes. In *Sabre Research Internal Technical Report*, August 14.
- Vinod, B., R.M. Ratliff, and V. Jayaram. 2018. An approach to offer management: Maximizing sales with fare products and ancillaries. *Journal of Revenue & Pricing Management* 17: 91–101.
- Zouaoui, F., and B.V. Rao. 2009. Dynamic pricing of opaque airline tickets. *Journal of Revenue and Pricing Management* 8 (2/3): 148–154.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

B. Vinod serves as Senior Vice President and Chief Scientist at Sabre. Before rejoining Sabre in 2004, he was Vice President of Sabre Airline Solutions, responsible for pricing and yield management.

