The Internet and increased competition among airlines is forcing carriers to redevelop their revenue management systems to ones that provide fare visibility, remove restrictions and have smaller fare increments. Start-up airlines must consider their revenue management options carefully.

RM considerations for developing airlines

Revenue management (RM) is a critical component of any airline's operation; airlines with effective RM processes consistently outperform their competitors. Increased competitive pressure, combined with passengers able to access cheaper fares, has made RM more critical than ever. Established airlines have to adapt their systems and processes to meet these changes. This can be difficult, and new or developing airlines are able to establish effective methods with greater ease.

RM is the practice of maximising revenue from the sale of airline seats by controlling price levels and seat availability to balance supply with demand. Airlines have questioned the validity of RM, due to its failure to arrest the decline in passenger yields and the loss of passengers to lower fares or low-fare airlines. RM is still valid, but the evolution of the passenger market has placed new demands on airlines.

The evolution of the aviation market includes decreased price levels, increased choice and aggressive competitors. These factors are forcing larger carriers to adopt the RM processes and strategies of their low-cost rivals'. Many network carriers, which once relied on complex fare rules and expensive ticket prices, have had to re-align their pricing and RM processes to reflect the new market conditions. The RM methods of larger carriers are changing to ensure competitive market positioning. Smaller carriers, or those evolving their business practices, need to examine the RM options available to them to ensure that the processes they develop are both practical and effective.

RM Evolution

RM, also called yield management, was established by American Airlines (AA) in the 1980s as a response to the threat from People Express (PE). AA needed a way to compete with PE's low fares, and realised that if it could sell seats that were flying empty it could match or beat its competitor's prices. The challenge was to stop passengers who were willing to pay more for a seat accessing a lower fare.

AA developed powerful mathematical algorithms to forecast future passenger demand for a flight, simultaneously imposing restrictions on cheap fares to discourage business travellers who required ticket flexibility. Combining ticket restrictions with the segregation of passenger demand into discrete sections based on price sensitivity has been the predominant RM model ever since.

This model of RM changed little until the late 1990s. Business passengers paid more to ensure they could travel when they wished; leisure travellers paid less while and accepted they could not change dates or departure times. The restrictions a passenger was willing to accept in exchange for a cheaper fare defined the passenger's price sensitivity threshold.

Forcing a passenger to pay as much as possible was possible through two primary methods: restrictive fare rules that eased as the ticket value increased; and the absence of low-fare seats on popular flights that coincided with business travel, thereby forcing low-revenue passengers on to emptier flights later in the day. This was an extremely effective process, which overcharged the business segment and inconvenienced the leisure segment. This allowed low-cost carriers to tap into the market, offering low fares on all flights, even early morning/late evening flights which cater primarily for the business market. The popularity of these carriers dramatically eroded other carriers' market-share, forcing them to match many of the low-fare policies and procedures. This has caused significant problems for many airlines, since their RM systems are not suited to the new operating environment without significant changes to RM procedures.

RM Functions

RM involves forecasting the future demand of each flight based on its historical performance. Accuracy requires large amount of data and statistical analysis. Traditional RM procedures were built on the premise of protecting seats for high-paying passengers, who generally book at the last minute. These passengers pay a premium for booking late, effectively reimbursing the airline for the risk incurred from not selling the seat earlier, albeit for less money.

The RM system forecasts demand in each booking class, and automatically adjusts the seats available in each booking class accordingly. Therefore, if a flight experienced high demand from business travellers, its lower fare classes would be closed, with only expensive classes available. Conversely, if a flight had little demand the RM system will make every fare class available to minimise the risk of having unsold seats.

Determining the classes that are available is based on calculating which solution provides the greatest revenue to airlines, based on risk versus reward. This process calculates the odds of selling a seat at high revenue versus the odds of selling a seat at lower revenue. This process is defined as Expected Marginal Seat Revenue (EMSR), and works as follows: An airline has one seat left for sale, which could be sold in either one of two classes. Class A's price is $1,000, while class B's price is $550. If there is only a 50% chance of the seat being sold in class A, the RM system calculates the seat's actual value at $500.

Class B, with a 100% chance of sale, would have a higher potential revenue value ($50 higher) and the system would assign the seat to Class B, even though the fare is lower than Class A.

EMSR is a complex process that calculates purchase risk on every seat for every class and assigns inventory

accordingly. When airline managers see strange inventory assignment, such as low revenue classes open on a very full flight, it is often an EMSR process based on the probability of selling the seat at a high fare.

Part of RM is for airlines to determine accurate EMSR boundaries. Many low-cost airlines evade this problem with smaller price increments per class, about $10-15. If EMSR calculations are inaccurate the revenue loss is not extreme. In the earlier example, if the RM forecast was wrong, the airline lost $450 in revenue, while a low-cost airline might lose $10-15.

Optimisation processes

The RM system is linked to the Global Distribution Systems (GDS), which display the airline's available seats (the inventory) to booking passengers.

Based on the prediction of remaining demand, the RM system closes lower fare classes from sale. This prevents passengers from buying cheaper fares, forcing them to pay more. As a general rule, the closer to departure time the booking is made the more it will cost. To further reduce this problem, airlines developed different fare classes in a hierarchical system. This places the full economy (Y) fare as the prime class and the other classes as subsets. Each class is able to access the seats for sale in its own, or lower, classes, but is unable to access seats in higher classes. The system allows all seats to be sold at the Y fare, while restricting lower classes to only those seats allocated to them by the RM system. People would book in different classes based on the level of restrictions for which they are willing to pay.

The management of inventory is the principal and critical function of RM. As bookings are made, the number of cheaper seats available should generally decrease. However, if the flight is not reviewed frequently enough this will not happen. Many systems only look at flights that are selling at a rate outside its normal booking profile during the early months of selling, usually from 12 months to 2 months before departure.

"Many airlines focus on up-selling, the process of forcing passengers to pay more for a seat, as the key role of RM," says Helmut Mokosch, senior vice president of network management with Germanwings. "This does not address the other issue, down-selling, where passengers pay less for a seat than they were willing to pay. These processes are opposite sides of the same coin; the need to manage available inventory smarter. Analysing booking periods shows that higher revenue purchases are generally made between 11:30 and 15:00, with the peak around 12:30; which coincides with people accessing the internet during lunch breaks. Another high revenue period is late Sunday evening, when people have been told to get to somewhere quickly on Monday morning. Lower revenue bookings generally come in during the evenings and the weekend. If airlines began to analyse their available data they would probably see similar trends. These trends could enable them to display high fares during the day and cheaper fares during the evening/weekend. Airlines are now trying to develop ways to manage passenger booking patterns now that fares are no longer segmented, but they have overlooked the point that passengers segment themselves."

RM trends

The pressure created by lower fares and an aggressive competitive environment has required RM to change focus. Many airlines were previously investing in sophisticated Origin and Destination (O&D) systems that sought to maximise revenue on as many flights as possible. O&D systems were designed to evaluate the benefit of a passenger willing to pay a higher fare, but only when flying one leg. For example, Paris-London, rather than the passenger wishing to pay a lower fare, but wanting to fly Paris-London-New York. Sophisticated mathematical algorithms were developed to evaluate the many revenue opportunities generated by an airline's network. O&D systems encouraged airlines to look at the revenue generation potential of their network, rather than focus on individual flights.

The increase in competitive pressure, especially in Europe and the US, is pushing airlines away from sophisticated O&D systems toward simpler solutions that are more suited toward price sensitive markets. "RM systems work very effectively when demand exceeds supply," says Alan James, senior manager of networks and revenue management with BMI British Midland. "They are less effective when the reverse occurs. RM systems will default to opening all fare classes to fill the flight to reduce the volume of unsold seats. This removes the potential to force passengers to pay higher fares. The new competitive environment increases the risk that airlines will be unable to force passengers to pay more because the passenger now has the choice of a competitor. The big

British Midland worked with Sabre Airline Solutions to adapt Sabre's Airmax system to meet British Midland's new market dynamics. British Midland now makes fares visible, offers one-way fares and has removed restrictions on its Internet booking site.

RM development areas consequently involve determining whether the competitors are offering higher or lower fares, and creating RM functionality that can manage pricing structures that have moved away from return to one-way prices. To develop effective RM you need to move away from inventory optimisation and develop price optimisation."

Price optimisation

With fare rules removed, the only control on pricing is now the availability of cheap fares. RM previously could display multiple fare options, but now it can only display one. RM has to determine the value of each aircraft seat based on the time before departure that a person wishes to purchase it. Booking well before departure enables a passenger to secure a cheaper seat than if they attempted to book the day before departure. Low-cost airlines also observe this rule; prices generally increase as the flight fills up the closer to departure.

Low-cost airlines generally have higher load factors than their competitors, which creates an interesting dynamic. Premium airlines are actually often cheaper than low-cost airlines, especially closer to departure, but the public perception is the opposite. An example of this is BA, which competes with BMI and easyJet on the London-Paris route. BA and BMI often offer one-way fares at about $30 for travel from Heathrow, while easyJet offers fares for about the same value, but departure is from Gatwick or Luton. Low-cost airlines can be more expensive than their competitors. Ryanair can demand one-way fares as high as $200, while full-service carriers charge less.

The low-cost airlines rely on demand to fill flights. As the flight fills up, lower fares are removed from the plane. This is different from network carriers, which attempt to segment the market via fare rule and then sell high- and low-revenue tickets simultaneously. In the low-cost airline pricing model there is no market segmentation. Fares are increased in steps, depending on how full the flight is. Low-cost airlines have small fare increments, of perhaps $10-20 per increment. The assumption is that a passenger may be prepared to pay the increment if it can secure them a seat. Other carriers are beginning to adopt this method.

Airlines are still grappling with price optimisation, even though it is relatively straightforward. The price at any point in the selling process will determine the product's attraction to customers. If too many low-value fares are available the flight will fill up quickly and late-booking, higher-revenue passengers will be turned away. Keeping the fares high risks the flight having empty seats that could have been sold. Airlines still have to fully master the relationship between demand and price. One airline that has tried a novel approach to pricing is Helvetic, a Switzerland-based start-up low-cost airline. Helvetic has only one booking class and price: $90. This means that the first and the last seat have exactly the same price, making no adjustment for demand. While this is a novel approach, it fails to maximise revenue opportunities and is therefore not very practical. Helvetic's main attempt with this model is to remove the need for RM, with the exception of managing overbooking.

RM systems

Large network carriers generally invest large amounts in sophisticated RM systems to manage the complexity that their networks generate. This improves their revenue by a minor amount; an example of the law of diminishing returns. Smaller carriers, particularly low-cost, invest in smaller and less sophisticated systems. The type of system required is a major part of the airline's business model.

Network carriers have had to re-design most of their RM processes and systems to accommodate the changed market conditions, prompting many to question the value of sophisticated O&D systems that have little application in the new airline environment.

BMI worked with Sabre Airline Solutions to adapt Sabre's Airmax system to meet the new market dynamics that BMI had to contend with. The removal of pricing structures to enable BMI to offer one-way fares required significant changes to the way BMI practised RM.

"We operate in a market that is intolerant of fare rules," says James. "We now operate in an environment where the only control on low fares is whether the fare is available when a passenger checks our website. Our new fare structure is all one-way pricing, which further reduces our options for imposing restrictions on passenger fares. The major change that we have had to implement is dynamic pricing. Pricing is no longer fixed; we adjust it to reflect booking demand. This means only the lowest fare is selling at any one point in time. Previously, it was possible to sell low fares and high fares simultaneously by relying on fare restrictions to limit sell-down. We needed to work closely with Sabre to adapt Airmax to new market conditions."

RM Options

Despite the complexity of RM, and the volume of data that it processes, it is wrong to assume that RM is a system or IT function. Two airlines may have the same RM system, but the system's effectiveness is driven by the ability and competence of the user, not by the forecasting algorithm it employs. One leading Asian airline believes nearly all the failures in its RM processes were primarily user-influenced; the users were not as well trained and motivated as its competitors.

RM systems do not have to be sophisticated forecasting platforms. A start-up airline, with perhaps 10 or less aircraft, could use general management tools, like Microsoft Excel, to perform RM. Forecasting demand based on historical data, or allocating inventory based on current booking demand, can both be performed in Excel, provided that the user is competent. Using established statistical procedures it is possible to determine the allocation of seats into individual booking classes with acceptable levels of accuracy.

Above this level there are many RM solutions available on the market, ranging from small consultancies offering tailor-made solutions to large vendors offering a standardised product. "The majority of network carriers will choose between PROS and Sabre solutions, as they are the market leaders," says Lucio Graziani, IT manager of the Volare Group. "Management are generally uncomfortable placing revenue responsibility into the hands of smaller companies, especially if they are an established carrier. Newer airlines are more adventurous, and are willing to use the skills of smaller vendors. Most of the systems employ similar methods for forecasting demand. The greater the complexity the longer time required for forecasting, however. Generally network carriers re-forecast each night. They begin at 10pm and finish around 5am. That method worked when people booked through travel agencies, but the internet has made that process redundant. With the internet providing 24-hour access, re-forecasting must now be virtually constant. Smaller systems can accommodate this change, but highly sophisticated ones are still catching up."

Low-cost airlines generally use Navitaire's Openskies system, which has a virtual monopoly on low-cost airline users. "Navitaire was the obvious option when we evaluated our options for a low-cost airline RM system," says Graziani. "The RM system we employed in our network airline, Volare Air Europe, was incapable of performing the same functions at the same price. Because it is less sophisticated, and does not have to calculate network origin and destination (O&D) possibilities, Open Skies can re-forecast flights constantly. Our previous system was deficient in this area, since it reviewed and analysed huge volumes of data and all flights. The new system only examines critical flights. Because of this, the volume of flights examined is smaller and the complexity is less; both of which improve speed. With this speed we can often re-forecast a flight between bookings. This has driven us toward a more dynamic approach to RM, where we can review the flight after each purchase to determine if it is optimised. Class availability is now the only constraint in RM. This is a critical development that enables us to maximise our revenue opportunities."

There are few RM vendors, because the investment in research and development (R&D) required are a major entry barrier. Below are the major vendors in the airline industry:

PROS

PROS has an 80% airline market share, including 17 of the top 25 airline carriers. Its system is one of the most sophisticated and effective. PROS invests heavily in R&D, contributing to its position as the leading RM solution provider. Its clients include: Air New Zealand, SAA and Southwest Airlines.

Sabre Airline Solutions

Sabre's Airmax system was one of the first RM solutions available, and has evolved considerably. Its newest version, which includes O&D control, is used by many airlines including BMI, Kuwait Airways and Qantas. It is unclear whether the developments made in conjunction with BMI will be made available to other airlines.

Lufthansa Systems (LSY)

LSY offers its ProfitLine O&D system, which has a unique O&D design based on booking record management. LSY estimates that a properly developed O&D system can add up to 2% of revenue for an airline. Lufthansa is LSY's major client.

Talus

Talus was the first company to develop an RM product after AA, and subsequently was an industry leader. Its position in recent years has been eroded by larger companies with greater R&D budgets. Talus has positioned itself primarily as an RM consultant which can also offer tailor-made solutions.

Navitaire

Navitaire's product, Openskies, is the de-facto system used by low-cost airlines. Developed in conjunction with Hewlett

Packard, Openskies is an effective solution that offers low-cost airlines great flexibility at minimal cost. The Openskies client list includes all low-cost airlines, except for easyJet.

Each vendor, as well as dedicated consulting companies like Veritec Solutions and The Rubicon Group, offers a broad range of solutions. The best solution will depend on the business model that it is using and the demands it has. All the systems available will perform capably and will manage the revenue enhancement opportunities.

Measuring RM benefit

RM has often been criticised because its value is hard to quantify. It is difficult to measure RM performance because passengers were not able to make a booking on a full flight, or the passenger who wished to pay less (or more) for a flight and was forced to pay a different price is not recorded, so the effect on future RM can be accounted for. How, therefore, can the effectiveness of an RM department be measured? One method, employed by many airlines, is the improvement in revenue per available seat-kilometre (RASK). RASK divides the total revenue generated by a flight over the available seat-kilometres (ASKs) that the flight generates. Improved RASK means that the revenue generated by a flight has increased. A continual improvement in RASK, adjusted for inflation, measured over a 12-18 month period indicates the effectiveness of RM.

Yield, calculated by dividing passenger revenue over distance, was previously used as the standard measurement tool. RASK and yield are not complementary methods, since RASK measures the total revenue per flight, while yield measures the revenue each passenger pays, and does not take load factor into account. Airlines may often block the availability of low fares on a flight in an effort to improve the yield. This results in the aircraft departing with empty seats and reduces RASK. Filling each seat will increase RASK, but depress yield. Airlines have often travelled with empty seats in an effort to protect yield, especially on business markets. This approach reduces the overall revenue potential per flight and is not an effective RM procedure. Yield should not be used for RM measurement.

Another performance measurement often used is spoilage: how many seats were eventually empty on a previously full flight? Spoilage indicates how well the RM system anticipates for no-shows, the people who book and pay for flights but do not turn up at the airport to fly. If a full flight is going to have a lot of no-shows, and therefore empty seats, the RM system will adjust for this by overselling the flight. Overselling involves selling more seats than the aircraft physically has. The measurement of spoilage and spillage are good indicators of the effectiveness of RM systems.

"It is extremely difficult to quantify RM's contribution," says Tim Claydon, senior vice president of sales and development for jetBlue. "It is a very contentious issue where agreement is hard to reach. The RM system may show declining revenue, indicating poor RM practices, but the market in general could be declining even faster. In this scenario the RM system has done a good job in maximising revenue in a shrinking market. Generally, RM measurement should factor in some form of annual pattern to enable the use of historical perspective. We evaluate our RM performance on year-on-year changes that are adjusted for seasonality. We also evaluate the RASK improvement that has occurred. If this is a constantly upward trend then we realise that RM is working. If it is erratic that indicates that seasonality and market forces are having a greater influence and we seek ways to address this. We use two main measurement parameters; we have found it is better to have one or two measurement methods, rather than multiple methods. Multiple measurement methods can create confusion, since they will generally be contradictory, primarily because they measure disparate issues."

Summary

Filling empty seats to improve revenue while stopping higher-paying passengers from buying cheaper fares is the core function of RM. The success of this depends on whether the airline has developed effective RM systems and practices. Developing practical RM processes takes effort and investment in tailored IT systems. Airlines need to know their requirements, what type of market conditions they will operate in and the competitive pressures they will face.

Competitive pressure from low-fare airlines requires a new approach to ensure parity is maintained. The only control mechanism now to force people to pay more for a seat is the timely closure of lower revenue buckets. This can be achieved by getting airlines to move to simplified, but effective, methods. AC