Passenger spill & spoilage is still a challenging element of revenue management. The problem originates from the difficulty in predicting demand for each fare class. Improved forecasting has led to reduced spill and made higher load factors possible.

Techniques to reduce passenger spill & improve load factors

Spill and spoilage of passenger revenue continues to be a problem for airlines. Airlines have managed to increase load factors by about 10% over the past 10-12 years, which suggests that they have managed to improve their controls of spill and spoilage. Modern revenue management (RM) and forecasting systems are better placed to reduce passenger spill and spoilage.

Nature of spill & spoilage

The simplest form of spill occurs when demand exceeds aircraft seat capacity. More specifically, spill relates to lost passenger demand and sales at a particular fare class. The revenue optimisation process in the period leading up to each flight calculates a bid-price curve at regular intervals. This uses several criteria to calculate a number of target seats, of the remaining available seats on a flight, to be made available at each fare or booking class. The bid-price curve calculation is a complex process, and uses several criteria. These include forecasting levels of passenger demand.

"Modern RM systems will also take into account price elasticity and the interaction between fare classes," says Hari Subramanian, director of revenue management at Sabre Airline Solutions. "Traditional RM systems assumed that the demand from potential buyers for each fare class was mutually exclusive. That is, the availability or non-availability of one fare had no effect on demand for the next lowest or highest fare. So, if a low fare class that had been available in the bid-price curve sold out, the demand for the next higher fare class was unaffected. The traditional RM systems also assumed that demand for each fare class conformed to a particular standard distribution curve.

"These assumptions are of course not true. When a lower fare class is closed off and the next higher one becomes available inevitably some potential buyers will not buy, but others will. The demand for the next available fare class therefore changes," continues Subramanian. "New generation RM systems appreciate that there is interaction between different fare classes. That is, if there are two fare classes available to a buyer, all buyers will select the cheaper fare class. This is in comparison to the traditional RM systems, which assumed there is no interaction between different fare classes when more than one are available to a buyer. Modern RM systems appreciate that this is not true.

"The reality of course is that as a lower fare class is closed off, only a portion of all potential buyers will buy the higher fare that becomes available," continues Subramanian. "This highlights the first problem of passenger spill, since it can never be known how many buyers and how much revenue is lost when a fare class is closed off and a higher one becomes available. New generation systems will try to calculate how many buyers are lost, and how many have been prepared to uprate to a higher fare. They therefore calculate the inferred price-elasticity. This is used as an element of future bid-price curve calculations."

By the time a flight closes the number of seats sold at each fare class will have been determined. Airlines clearly aim to maximise the number of seats sold, and the number sold at the highest fare classes. "If the bid-price curve estimation process used a demand forecast that overestimated the demand for low fare classes then too many would have been made available at the low fare class, and too few at higher fare classes. Lower than expected demand for low fares or higher than expected demand for high fares mean that more seats could have been sold at high fares, and so revenue is lost. This type of loss is known as spoilage," says Jean-Michel Sauvage, head of Altea RM system at Amadeus.

"If the bid-price curve estimation process has overestimated demand for high fares, too few seats would have been made available at low fares and there would have been insufficient demand to fill all the remaining seats available at high fares. This would result in some seats going unsold, when they could have been sold at low fares. Passengers that would have bought these seats were rejected and went to alternative carriers. This is known as spill," continues Sauvage.

Revenue optimisation

The challenge of calculating the bid-price curve is that demand varies, and there is a distribution of demand for each fare or booking class. "Perfect forecasting means that the exact level of demand for each fare class can be predicted," says Subrahanian. "In fact it is impossible to know how many buyers are spilt each time a lower fare class is closed off."

When the bid-price curve has been calculated for the day, there are several fares available for the same date and time of travel, and each would have different rules and restrictions. These different fare

One feature of modern revenue management systems is their capability to calculate the elasticity of demand for each fare class, and so the demand for a higher fare class when a lower fare class is closed. This improves the accuracy of the bid-price curve calculation. Not only has this reduced spill and spoilage, but also improved load factors and yields.

classes can be seen by a buyer when requesting a fare on an airline website, after they have input their desired date and time of travel.

The revenue optimisation process aims to calculate the marginal trade-off. The point of marginal trade-off is where there is a point of indifference between a lower and higher fare class. This is where the airline is indifferent to whether it sells lower or higher fare classes, because the final revenue will be the same. This means that if the bid-price and inventory control system is set at the point of indifference, then the airline will generate the same amount of revenue at whatever fare is set.

Subrahanian explains that, optimally, an airline wants to have the right number of seats available at each fare so that there is no wastage of unsold seats, or the problem of passengers being unable to get a particular fare because it has already sold out. This way the number of seats made available at each fare ensures a balance between spill and spoilage. This is because the number of seats made available at each fare class equals demand for seats at each fare class.

Demand forecasting

Revenue optimisation, and minimising spill, ultimately depend on high quality forecasting. "When forecasting is perfect then the forecast demand for each fare class equals the actual demand that materialises," says Sauvage. "This would then allow the number of seats set in the bid-price curve to be to be calculated perfectly. This is actually impossible to achieve."

Several hundred factors affect actual demand for each fare class. Not only are these too numerous and complex to model, but they are also too hard to monitor. This includes the fares and available fares offered by competitors. While some can be monitored by web crawler software, not all fares offered by all competitors through all sales channels can be accurately tracked.

There are also a large number of economic factors external to the airline's influence that vary on a continuous basis.

"A main difficulty is the large number of factors that affect a buyer's behaviour," comments Subrahanian. "A major objective of demand forecasting is to be able to model a buyer's purchasing behaviour.

"Another major objective is making an accurate assessment of the interaction between different price points or fare classes," continues Subrahanian.

"Improved forecasting also includes accurately predicting variations in demand that occur over different dates for the same flight number and departure times; variations in demand for the same flights between two cities during the same day, week, month and season; and variations in demand during and around special events," says Subrahanian. "The dates of special events, such as Easter and Thanksgiving, change each year, but there are also one-off special events. These factors have to be taken into account.

"Another element of improved forecasting is good intelligence about competitors' fares," adds Subrahanian. "As described, web crawler software has helped here, but this is insufficient to provide all the information required."

The importance of forecasting to minimise spill is illustrated by Sauvage's observation that spill increases with uncertainty or inaccuracy of demand forecasting. "It is also important to have a precise revenue optimisation system in addition to accurate forecasting," says Sauvage. "Forecasting is more important than optimisation because optimisation first requires a good feed of high quality forecast data."

An important element of forecasting is the feedback and use of historical data for future forecasting purposes. The information used in this feedback process is historical sales data, sales achieved to date for future flights, sales achieved since the last optimisation, the remaining seats available, the fares and fare classes available in the system, the price elasticity between each fare class, intelligence of competitor's fares, and how fast the flight is selling and the nature of the demand compared to previous flights.

A further refinement to revenue optimisation and reducing passenger spill would be the use of dynamic revenue optimisation. This is a recently developed technique where an airline's entire network is re-optimised after each booking is made, rather than once every 24 hours during the night. Dynamic revenue optimisation has only recently become possible and is expensive. It would, however, provide a higher degree

The advent of low-cost carriers over the past decade has led to a stimulation of demand. LCCs have nevertheless created overcapacity on many routes. Incumbent airlines have consequently reduced capacity through the use of smaller types while maintaining service frequencies.

of fine tuning and so reduce levels of spill and spoilage.

Over-booking policy

Another cause of spill is the loss of revenue due to passenger no-shows and denied boarding. Airlines are clearly aware that there will always be some passenger no-shows, particularly those holding tickets with high fare classes that are fully flexible, or cancellations in the last two weeks prior to a flight departing. There will also be some passengers that are denied boarding at the airport, often because of missed connections, fewer than forecast no-shows, over-aggressive over-booking policy, or poor forecasting.

These cancellations, no-shows and denied boardings all have to be planned for, and built into the optimisation system. "Airlines will therefore need to have an over-booking policy to account for these losses," says Sauvage. "The typical system is to allow up to 110% of a flight's seat capacity to be sold up to 14 days prior to departure, and then reduce this in stages. This can be down to 105% up to seven days prior to departure, and then at 100% by the day of departure.

"What actual degree of overbooking is allowed and on which days prior to departure all depends again on accurate forecasting," continues Sauvage. "This requires good quality historical data of cancellations, no-shows and denied boardings. The ultimate objective is to achieve a 100% load factor at flight closure, once all cancellations and no-shows have taken place, and without any denied boardings.

"The revenue optimisation system also needs to take account of predicted cancellations and no-shows at particular fare classes when calculating the bid-price curve. Amadeus is working to integrate the check-in and RM systems to better estimate the cost to an airline of a denied boarding," says Sauvage.

Load factor & spill

The problem of passenger spill is generally greater with higher load factors. That is, if demand is particularly low then many seats will be left unsold, even when only the lowest fare classes are available. This means there is unlikely to be any lost demand and spill. This can be overcome by reducing capacity through scheduling fewer flights or changing to smaller aircraft types.

Spill will generally increase as load factors and demand rise. Higher demand will result in a larger portion of seats being sold in high fare classes, and also higher overall revenues and yield mixes. It becomes harder to calculate the optimised bid-price curve as load factors get higher. This will inevitably lead to spill and spoilage in the revenue optimisation process. Forecasting and revenue optimisation therefore become more important as load factors get higher.

"Airlines have generally been able to achieve higher load factors over the past 10-12 years," says Sauvage. "This is due to better and more accurate forecasting and improved revenue optimisation. As a result, airlines are able to achieve lower rates of: spill, lost passengers, spoilage, and lost higher-yielding passengers. These provide more accurate predictions of demand at each fare class, and so a more optimised bid-price curve. More seats are therefore sold at each fare class made available in the bid-price curve.

"This ultimately leads to a higher yield mix, in addition to higher load factors," continues Sauvage. "Another change for airlines over the past 10-12 years has been an improvement in the distribution and sales channels they use. A wider range of sales channels means that airlines are able to offer a wider range of fares. They can therefore appeal to a wider range of buyers and generate a more specific bid-price curve in their revenue optimisation processes.

"More frequent revenue optimisations have also helped to improve yield mixes, increase load factors and lower spill and spoilage rates," adds Sauvage. "This is because frequent assessment of the sales achieved since the previous revenue optimisation process allows a more accurate assessment of demand and the price-elasticity for each fare class, and so ultimately a more accurate prediction of demand and the bid-price curve. More frequent revenue optimisation has become possible with the more powerful RM systems that have become available over the past decade."

The advent and increased market presence of low-cost carriers (LCCs) had generally led to over-capacity on many routes. While LCCs offer low fares that stimulate additional passenger demand, the LCCs have generally disproportionately increased capacity on many routes, leading to a dilution of fares and revenues for the incumbents. The response by the incumbents has generally been to reduce capacity through use of smaller aircraft types while maintaining service frequencies. AC