



WCTR

SPILL MODELING FOR AIRLINES

WILLIAM M. SWAN

Boeing Commercial Airplanes
Airline Industry Analysis
PO Box 3707 MC 21-28
Seattle WA 98124-2207 USA

Abstract

Spill models estimate average passenger loads when demand occasionally exceeds capacity. Such models have been in use for over 20 years. The shape of the distribution of demand is discussed from both theory and observation. Sources of variance are identified and calibrated. Measurement problems and techniques are discussed. Two alternate spill formulas are presented. A model revision responds to changes in process caused by computer reservations systems and revenue management. The concept that spill losses should be valued at discount fares is discussed. The recapture of spilled demand is presented as well as when such a phenomenon is relevant. Comparison of various sources of error is included. Finally, the use of spill models “in reverse” to imply demand from load is shown to have poor accuracy. The paper is meant to offer to the literature a reference for basic use. It is the result of 20 years’ involvement in spill model derivations, calibrations, and applications.

INTRODUCTION

Spill is the average passengers per departure lost off a group of flights because demand sometimes exceeds capacity. A group of flights can involve a flight leg for a month, season, or year. Groupings can involve one leg, a small group of legs reassigned from one fleet type to another, or all the legs served by a single aircraft type in a fleet. The spill model has been used widely for over 20 years within the airline industry. However, there has been no commonly available publication discussing its use. This paper attempts to put into the record the formulation, its calibration, and some issues of use for airline analyses.

The basic idea behind spill is that demand for a group of flights can be represented as a distribution about a mean. The integral of this distribution is the “fill” rate for seats on an aircraft. This is shown in Figure 1. The integral of the fill rate beyond a truncating capacity is the spilled demand. While spill is the term usually calculated, the model is commonly employed to estimate the difference in spill between two capacities. This is the fill rate for the extra seats. Perhaps the “spill” model should have been named the “fill” model. Time and tradition prohibit this nomenclature, but spill model performance is judged by its performance in estimating fill.

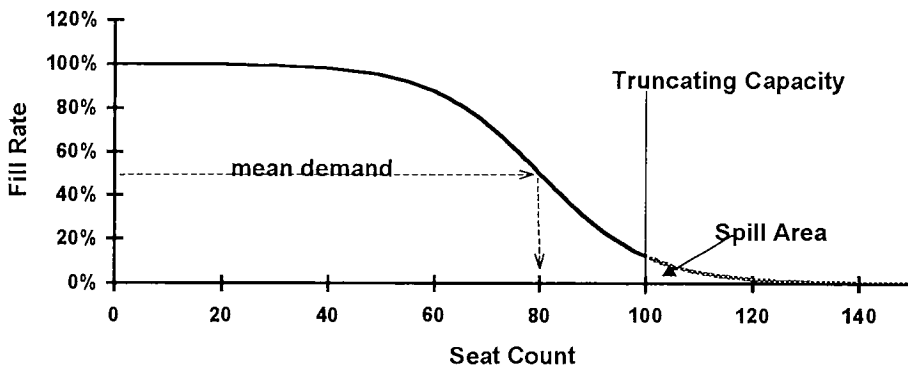


Figure 1: Fill Rate and Spill

The discussion below begins with the characterization of the demand distribution. When is it Normal and why or why not? Development continues with two formulas for calculating spill when the distribution can be thought of as Normal. Discussion then maintains that the appropriate fares to apply when valuing spill are discount fares. Arguments are put forth that the complications of the recapture of spill often can be avoided. Finally, discussion discourages the use of the spill “in reverse” to estimate demands from observed loads.

UNDERLYING STATISTICAL MODELS FOR DEMAND DISTRIBUTIONS

The idea behind spill is that the demand for a series of departures has a distribution about its mean. The most central case in the airline business is the distribution of demand for a flight over a month.

For example, the demand distribution for the set of 30 executions of the 9:00 flight from Seattle to Chicago during the month of April. The amount of variation in the distribution can be measured by the ratio of the standard deviation to the mean. The convention within the airline industry is to refer to this ratio as the “K-factor” [DeSylva, 1976].

One source of the variation in demand is pure randomness. Imagine that 1 million people in Seattle are candidates to fly to Chicago at 9:00 on the first day of the month. Each person flips a coin with a one in 10,000 chance of coming up heads. The probability distribution for the number of heads from 1 million such Bernoulli trials would have a mean of 100, a K-factor near 0.10, and the shape of the particular Gamma which has the appropriate moments. This would look like the demand distribution for the flight for one day. With demands the size of an airplane, the Gamma shape is almost Normal.

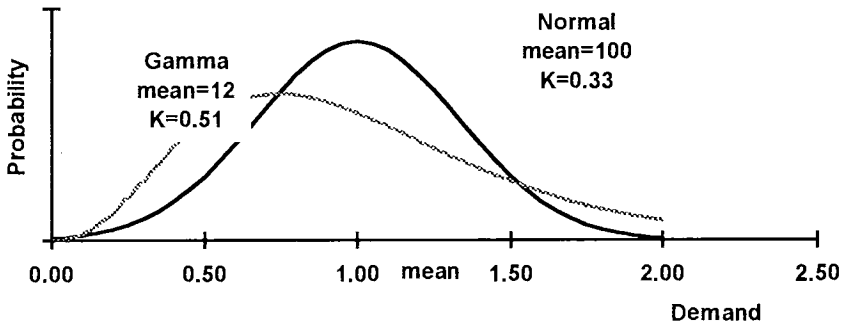


Figure 2: Normal and Gamma Shapes

For smaller demands the Gamma shape is more skewed and wider, as show in Figure 2. The small demand case is appropriate for first class demand, or for any small sub-group of fares or destinations on a flight leg. A one-day spill model as above should apply for revenue management planning, which focuses on such sub-groups.

The spill model for a month’s execution of this flight includes further variations. These further variations overwhelm the random effects, unless the demand mean is small. These variations come from the cycles of demand through the days of the week and the weeks of the month. Such cyclic variations correspond to changing the probability of heads for the coin-flip experiment day-by-day. It would be higher on Friday, and lower on Wednesday. This changes the expected demand day by day. For a single flight for a month, cyclic variations alone are enough to create a K-factor of 0.30. When cyclic variations are from several sources, or when they capture uncertainty in the estimate of mean demand, they are most naturally Normally distributed. Total variation becomes a combination of cyclic and random sources. The combined shape is a compromise between the Normal and Gamma with the Normal the dominant whenever the cyclic component of variance is greater than the random component.

Further cyclic variation occurs when considering not one flight leg but all the legs assigned to an aircraft each day. Still higher values apply for all the legs flown by one fleet type for the month, or by a fleet type over the 12 months of a year.

Proprietary data on demand distributions has been reviewed covering a large number of cases. Data has come from U.S., European, and Asian airlines covering both domestic and international flying. Data from 15 years back has been examined, as well as data nearly current. Most data has been daily onboard loads, but analysis has also been done on reservation system bookings. Problems with the data are discussed below, but the overall conclusions seem to be supported by most cases.

Table 1: Typical K-cyclic Values

case	Day	Month	Season	Year
Flight Leg	0.00	0.30	0.32	0.36
Aircraft Increment	0.18	0.35	0.37	0.40
Fleet	0.32	0.44	0.45	0.48

The primary conclusion is that the demand distribution is as close to Normal as anything else. Considering the multiple sources of variation, the central limit theorem would lead us to expect this outcome. This rule holds for demand means above 70. For smaller demands, Gamma, log-normal, or truncated Normal distributions are of interest.

The second most important conclusion is that K-factors are surprisingly constant across widely differing market types. Furthermore, the increments of variance seem statistically independent. Table 1 presents the results of studies of the cyclical components of K-factor. The random component is reserved for the subsequent paragraph. Table 1 says the day-of-month variation produces a K-cyclic of 0.30. Broadening this to the flights that would be assigned an incremental aircraft raises the value to 0.35. For instance, this would be the variation among 4 legs transferred from service by small aircraft when one additional large aircraft becomes available. This variation is driven by changes in demand by time-of-day. The K-factor for an entire fleet adds the spectrum of demands a fleet type is expected to serve. Usual circumstances would see a rise to 0.44 for this effect. The variations across the months of a year would drive this K-factor to a total of 0.48. Finally, planning studies often accept an additional uncertainty in the forecast of the mean demand of 20% or more. This can bring the total cyclic K-factor up to 0.52. This addition is not presented in Table 1, since it is not an observed variation.

Random variations add very little to these cyclical components, for demands of 100 or so. If everyone traveled alone, the standard deviation would be roughly the square root of the demand. However, the (root mean square) average group size is closer to 2, so the standard deviation is the square root of twice the demand. This increases a K-cyclic of 0.30 to a total K-factor of 0.33 for a demand mean of 100.

For small demands such as first class, the random variance is large and the demand distributions are not Normal. For demands below 3, the monthly K-factor can be above 1.00 and the shape can approach a simple exponential distribution. Some airlines and researchers employ low-order Gammas or other distributions for such cases. Both data and theory suggest they should. Oddly enough this approach is seldom followed in revenue management formulas, where it would seem to be most appropriate.

Increases caused by random variations are confirmed by looking at total K-factors for monthly cases. Monthly total K-factors in Figure 3 illustrate a decline from high values at low demands toward the K-cyclic asymptote at high demand. Figure 3 can be reproduced using detailed flight leg data.

Direct calculation of variance is difficult. Even with perfectly clean data, a month's worth of data points gives a poor estimate. Unfortunately, the data are far from clean. Loads and bookings are

truncated by capacity. Low loads are often the result of flights with delays or weather complications. High loads are sometimes the result of the cancellation of some near-by flight. These distortions focus on the tails of the distribution. Unfortunately, the tails of the distribution would provide much of the information about the size of variations, if the data were clean. Observed load variations consistently underestimate the variation of underlying demand. More careful calibrations almost always lead to higher estimates of K-factors.

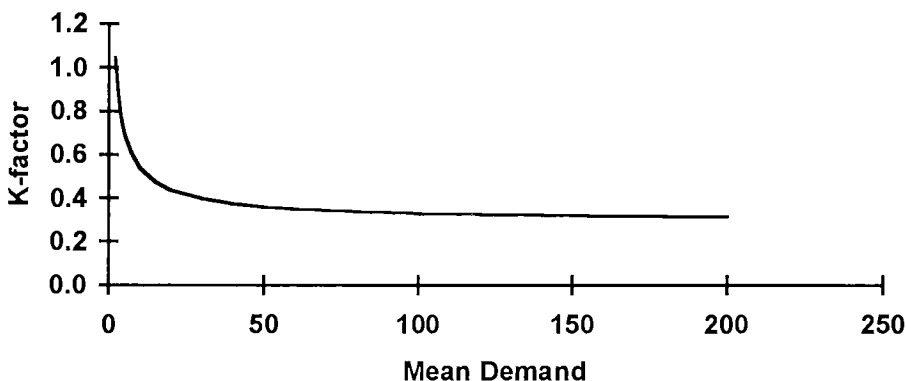


Figure 3: K-factor Rises for Smaller Demands

Practical calibrations of K-factor use the median to approximate the mean, and the distance from the median to the 25%ile observation to estimate the standard deviation. This gives up about half the formal statistical efficiency, but produces better results on real data. This has been tested by simulating clean data and simulating the usual distortions from truncation and delayed or canceled flights. The simulated distributions closely resemble real data. However, for simulations the “true” underlying K-factors are known. Estimates using the 25%ile and 50%ile loads capture the K-factors underlying simulations over useful ranges of **K**. A practical fit that works using the standard spill formulas even for smaller first class cabins is:

$$\text{K-factor} = (\text{Load}_{50\%ile} - \text{Load}_{25\%ile} + 1) / (0.674 * \text{Load}_{50\%ile}) \quad (1)$$

Calibrations of K-factor are best done in months with low load factors. Averages over several months are needed, even for clean data. Under few circumstances can the K-factor for an individual flight leg be estimated accurately. However, similar markets have similar K-factors, and values seem to be constant across a surprisingly large range of market types. Where data is unavailable, the values from Table 1 are often used.

K-factors for markets that are purely local and purely one kind of traffic are up to 20% higher than indicated in Table 1. This is because both business and pleasure markets have higher K-cyclic components than their combination. Not only are they somewhat independent, there is a small negative correlation between the two types of demand. They can be the same people taking different kinds of trips. Most data used for calibrations comes from flight legs with a mix of business and pleasure travel, and a mix of local demand and demand connecting beyond the local city-pair. Over a broad range of mixes, the common K-factors of Table 1 result.

SPILL FORMULAS

Presentations of the straight Normal version of the spill model formulas date at least as early as 1975 [Shlifer and Vardi, DeSylva]. $P(x)$ is defined as the probability P of demand x . $P(x)$ is Normally distributed with mean μ and standard deviation σ . $P(x) = N(\mu, \sigma; x)$. For truncating capacity C , the number of spilled passengers is $(x-C)$ and the total spill $S(C)$ is

$$S(C) = \int_C^{\infty} (x - C) \cdot P(x) dx = \sigma \cdot N(0, 1; B) - \sigma \cdot B \cdot (1 - \Phi(0, 1; B)) \quad (2)$$

Where $\Phi(0, 1; B)$ is the cumulative Normal, and $B = (C - \mu) / \sigma$.

This formula proved awkward in practice. It represented a small difference of two larger numbers and required accuracy in calculating N and Φ . There was no explicit formula for Φ , so a 5- or 7-term approximation had to be used. This made the formula difficult for early spreadsheets and relegated spill calculations to table lookups or use within larger scientific language programs. Modern spreadsheets contain functions for Φ .

In 1980 a simplification was made using the common logit approximation of the cumulative Normal [Swan]. This was not accurate enough for Φ in calculations using (2), but it allowed an alternative derivation. $F(s)$ was defined as the fill rate for seat s . The fill rate was the probability that demand equaled or exceed s . For $b = (s - \mu) / \sigma$,

$$F(s) = 1 / (1 + \exp(1.7 \cdot b)) \quad (3)$$

With $ds = \sigma db$, the integral of the fill rate for all seats above capacity C gave the spill value:

$$S(C) = \sigma \int_B^{\infty} F(b) db = (\sigma / 1.7) \cdot \text{Ln}(1 + \exp(-1.7 \cdot B)) \quad (4)$$

A further extension provided the displacement rate D . Displacement is the incremental spill for an addition of one customer a day to the average demand, $\partial S / \partial \mu$:

$$D = S / \mu + (C / \mu) * F \quad (5)$$

Displacement values are higher than fill values because an added customer is more likely to show up on a peak flight, while an added seat is added equally on all flights.

The simpler logit formulation meant spill was coded into spreadsheets and pocket calculators of the day. This popularized use in business practice. Use now ranges from aircraft assignments to a schedule for a month, to studies of seating configurations, to the costs of marketing promotions, and most critically to fleet planning. Most major North American airlines employ this formulation, as well as several major carriers in Europe and Asia. Other airlines maintain equivalent formulations using Normal or Gamma distributions. Log-Normal and truncated normal distributions can also be handled now with modern spreadsheet functions. Comparison of these various derivations is beyond the bounds of this discussion. It represents much-needed research.

Spill can be calculated numerically using any reasonable distribution as the underlying description of variations in demand. Earlier discussion of demand distributions suggests that the Gamma may

be the most appropriate for small demands and small groups of flight legs while the Normal may be best for broader applications. In any case, decisions are almost always based on the difference of spill between two capacities. This is the fill rate for the incremental seats. For typical demands, numerical differences between Gamma, Normal, Log Normal, and Logit versions compared by incremental fill values are small. For airline applications, such differences are overwhelmed by uncertainty in the estimate of the mean demand, K-factor, or other parameters. These estimates themselves are often best based on fill rate observations. The uncertainties of estimates will be discussed later in this paper.

REVISIONS

K-factors received a modest modification in treatment in 1983 [Swan]. Before then, K-factors were treated as independent of demand size. This implied that all variation was driven by cyclic factors. The random component was neglected in both discussion and estimation. A single K-factor for all fleet planning applications allowed the spill model to be a table lookup based on demand factor (demand divided by capacity) alone.

Recognition that there was a random component to variations explained some of the differences between very large and small aircraft and between total demand and demand for smaller component cabins. Revised versions of the spill model employ K-factors including both cyclic and random components, as already discussed. Cyclic variations do not depend on the size of the demand, but only on the case being studied. Random variations do not depend on the case, but are specific to the value of the mean demand. Overall variance is the sum of the two effects. This means that K-factors change slightly with demand. This is illustrated in Figure 3. For demand levels above 100, the random component of K-factor has been a complication with little numerical significance. For smaller demands, it has improved estimates meaningfully.

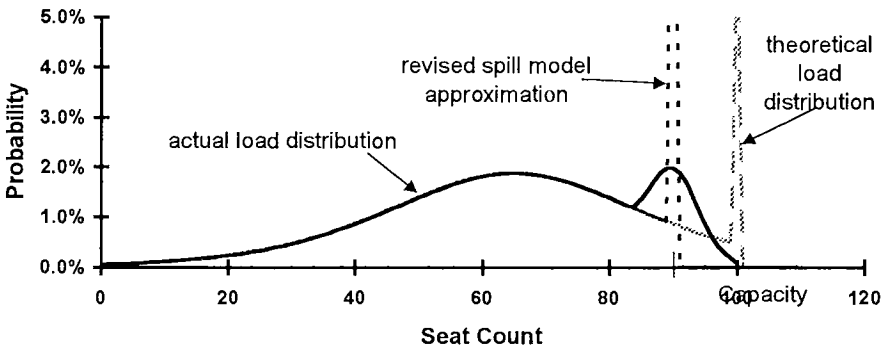


Figure 4: Loads: Theory and Practice

A second revision of the spill model changed spill values significantly. It was recognized that a flight’s “truncating capacity” is not the physical seat count. A flight is not full at 100% load factor. It is full when reservations are no longer accepted. The limited number of reservations translates through no-show behavior to a load at the gate. Optimal overbooking policies [Schlifer and Vardi] mean that the expected load is 5%-10% below the aircraft seat count. This 5%-10% is called “spoilage” in airline parlance. Spoilage averaged below 5% in the days of a single fare and reliable

no-show behavior. In those times spoilage served solely to protect against excess overbooking, preventing denied boardings at the gate. With discount pricing and revenue management there is a second reason for spoilage. Revenue management holds some seats open for late-booking high-fare demand. This demand does not always materialize, but airlines are willing to take the chance, since revenues run three times the discount fares. When these seats are not called for, they add to spoilage. With discounting and revenue management, the average truncating capacity dropped toward 10% below seat counts.

These issues are illustrated in Figure 4. The “theoretical load distribution” has an impulse function representing the 100% load factor cases of flights being full. This is the old spill approach. The “actual load distribution” has a small hump of load outcomes in the 85%-95% load factor range that represents the loads at the gate for cases when discount reservations were no longer accepted. Simulations have shown that for spill calculations it is sufficient to represent this “hump” as an impulse function at the new truncating capacity, in this example at about 90% load factor.

The spill model was modified to use the new lower capacities. Needless to say this increased the estimates of spill. However, the understanding that spill was largely discount demand decreased the cost of spill.

Optimal spoilage levels involve an interaction of the overbooking and space-protecting aspects of revenue management. Simulation shows that the optimal spoilage levels for a given mix of fares and uncertainty rises with the square root of the aircraft capacity. That is for seat count R , spoilage s would be:

$$s \approx c \cdot \sqrt{R} \tag{6}$$

Studies showed that the factor c should be as low as 0.5 for a single fare case and could rise above 1.0 for cases with discounts similar to current US conditions. Large c values were appropriate when discount fares were low compared to full fares or when uncertainty in no-show rates was high.

Average spoilage for an airline implies the value for c . Spoilage should be deduced from the departing loads of flights that are closed to discount fares. When a flight is closed to discount fares, demand is being spilled. It is not appropriate to measure spoilage only from flights that are closed to high-fare levels. These have low spoilage, but they are not representative. Nor are flights that are almost always full typical. They often record unusually low spoilage. This could be because no-show rates are more predictable, or it could be because stand-by demand tops up the loads. For most planning uses, it is appropriate to calibrate spoilage from a broad representation of closed flights and not from only from flights that are closed to full fare or closed frequently. While it is important to include spoilage in predictions of fill or spill, results are not overly sensitive to getting the exact spoilage value correct.

The concept of spoilage changes the spill formulas in a simple way. The capacity parameter C becomes the seat count less the spoilage:

$$C = R - s \tag{7}$$

REVENUES FOR SPILL

The spill model predicts spilled demand. The natural question is, what is the revenue for that spilled demand? The discussion of spoilage recognized that spill takes place by turning away discount demand as it requests a reservation. Revenue management systems' function is to spill discount demand and maintain space for higher fare demand. So the question of what revenue is spilled is either very complicated or completely simple. The complicated answer involves understanding what a revenue management system is trying to do on a detailed level, and how well it succeeds. The simple answer is that spill is at the local market discount fare. Discussion will try to motivate the simple answer.

The purpose of a revenue management is to spill discount fares when spill must occur at all. Most current revenue management systems group fares in to "buckets" and limit sales from the lowest fare bucket. A typical flight leg is half local traffic, and the local traffic is usually well over half at discount fares. Local discount fares are lower than connecting discount fares. So most revenue management systems limit local discounts first.

Even the best revenue management systems do a poor job of spilling just discount when load factors are low and spill is small. However, when significant numbers of passengers need to be turned away, it is easier to deny mostly discount demand. Furthermore, spill applications value differences in spill. That means it is not the average fare turned away that counts, but the expected fare for one last increment of spilled demand that counts. Simulations of leg-based revenue management systems suggest that when spill is not too small, 80% of it is turned away at the discount fare, and only 20% at an average mix of fares. This split is fairly consistent from modest levels of spill up to very high levels of spill and over a range of discount market shares and prices. The rule breaks down at high levels of spill, when all the discount demand has been denied and higher fares need to be refused. The practical conclusion is that spill revenues are just above the discount levels.

The most advanced revenue management systems try to do better. Origin-Destination based systems try to turn away discount demand from two-leg connecting itineraries if both legs are likely to be spilling. The revenue lost per leg becomes only a share of the connecting discount fare. The value is well below the local discount fare. This line of reasoning means average spill is at revenues slightly below the local discount, not slightly above.

Overall, spill is at the local discount fare, or at a value within 10% of this number for planning cases. This is well within the uncertainty of estimates for other parts of a plan, For markets such as domestic US hub services, the value runs about 75% of the average yield allocated to a flight leg.

RECAPTURE OF SPILL

Recapture is the idea that spilled demand does not fail to take the trip. Some of it finds its way back on to other flights by the same airline. This is easy to visualize on a daily or weekly basis. Spill from the 9:00 flight will divert to seats on the 11:00 flight, and spill on the Tuesday departure can arrange to go on Wednesday. For a day or a week, spill applications need to address the issue of recapture.

Recapture is less of an issue for fleet planning. In fleet planning, spill in August cannot be expected to use space in February, and spill to London does not board the flight to Miami. While some spill does find space on adjacent flights, the last incremental units of spill are left with fewer

and fewer open alternatives. The broader or longer-run the case or the higher the spill values, the smaller the likelihood of practical recapture.

For the shorter-run, there remains a need for understanding recapture behavior. Recapture has been studied using demand models that simulate passenger choices and preserve the second, third, and fourth choice departures for spilled passengers. It is important to preserve a list of alternatives, since if a passenger has been spilled off his first choice, he is late booking. Other flights are likely to be full with their primary demand or earlier recaptured demand. While a list is important, preserving a long list presents a problem. At some point customers give up and replan their trip around a different set of times or days. Nonetheless, the list-of-choices logic has been used in simulations covering a month of flights with day-of-week and time-of-day cycles. The results suggest the following simplification of recapture behavior: spilled demand for a city pair loads itself on flights as if it is seeking empty seats with little attention to schedule. After a first pass of primary demand and primary spill, the pool of spilled demand distributes itself at equal load factors on the remaining available space in the market. Available space is measured as the seats between the first-pass load and the truncating capacity C.

This result is a lot less certain than earlier statements about spill. Modest load factors and small spill will produce the more intuitive result that the more popular flights get most of the recapture. High demands produce the obvious answer that all available capacity is used, and the excess demand is lost entirely.

This understanding of recapture has an significant corollary. Extra seats on flights are not only useful for preventing spill from the flight, they also have value for accommodating spill from competitors flights or off other flights of the same airline. The reverse side of the "recapture" coin is this constructive use of extra capacity. The term suggested for this phenomenon is "refill."

Recapture means spill is less costly than it seems. This means extra seats are less valuable. Refill means extra seats have increased value. The two parts of the recapture phenomenon do not exactly cancel out, but they can be of similar size. It is not correct to include one without the other.

In annual or fleet cases recapture is small, particularly for incremental changes. For monthly cases for a single flight leg, recapture can be important, but the phenomenon of refill cancels some of its value. Overall, recapture requires a great increase in complication. Unfortunately, is also is an area where current research has leaves a great deal of uncertainty. Many applications choose to argue that recapture and refill are second-order effects and leave them out.

ERRORS IN ESTIMATION

Spill can be estimated, but how good is the estimate? Spill calculations require estimates of a number of parameters. The mean demand, K-factor, spoilage, average revenues, recapture, and refill all have uncertainties in their estimates. The way to test these estimates is to compare their effect on the value of an incremental seat on the aircraft. The use of a seat is its fill rate. This percentage will then be multiplied by the expected fare for spilled passengers to get the value for an extra seat. The table below develops an example with variations reflecting the separate uncertainties. Estimates of the uncertainty of individual parameters is from experience in spill applications.

Table 2: Errors in Spill Value for a Flight-Leg Month

Parameter	Estimated Value	Estimated Uncertainty	Range of Result (Seat Value)
Capacity (base \$)	200	±0	100%=\$31
Demand	150	±15	55%-155%
K-cyclic	0.30	±0.05	84%-116%
spoilage factor "c"	0.85	±0.15	94%-106%
Spill Fare	\$150	±\$25	84%-116%

Table 2 shows a value of \$31 for a typical case of a spill for a flight leg for a month. This represents using spill to help decide which aircraft type to assign to flight legs in a published schedule 3 months before the schedule will be flown. The \$31 is the value of an incremental seat using the estimated values for the list of parameters. Demand uncertainty has the largest range of values, from \$17 to \$48. Uncertainty in the mean demand reflects not so much the forecast of total industry demand three months ahead as the uncertainty of allocation for one particular flight leg month after month. Uncertainties in the fare represent both the difficulty of forecast for a single flight leg and the controversy about exactly which fare is spilled. Uncertainty in spoilage has only 10% of the range of demand uncertainty, and doubt about the proper estimate for K-cyclic is under a quarter. For such a short-run study, recapture and refill values would be relevant. Uncertainties in these would equal demand effects.

For fleet planning, fill values are higher for the same demand factor, because the data for a fleet for a year has more variation from its average than a flight leg for a month. Although fill values are higher, the errors are lower. Averaging across an entire system reduces the uncertainty in K-cyclic and spill fare estimates, and recapture and refill are much smaller issues for annual and fleet spill.

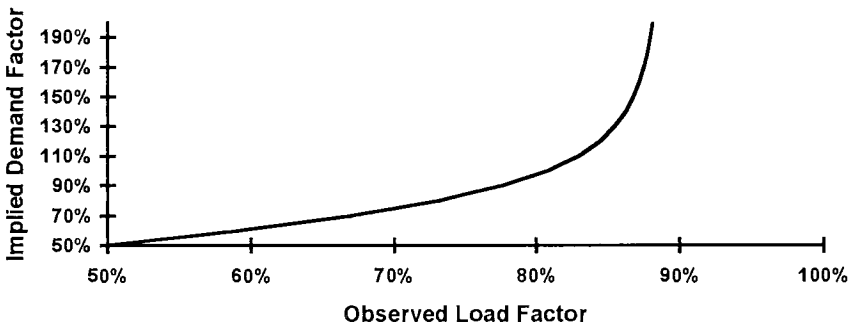


Figure 5: Implied Demand Skyrockets at High Load Factors

ESTIMATING DEMAND

The spill model starts with an estimate of the unconstrained demand for a flight. Often, this estimate comes from historical loads. This is fine when load factors are low. An iterative process

can establish what the demand would result in the observed load. However, when spill is an issue, load factors are already high. With high load factors, most flights are full and there is little information in the load distribution. It is very hard to determine what the underlying demand distribution was from the shape of the observed load curve. Another way to see this is shown in Figure 5. Implied demand factor is shown against observed load factor for a flight leg for a month. This is based on numerical inversion of the spill formula. Above 85% load factor, as little as a 0.5% point rise in observed load implies a huge increase in demand. To make matters worse, errors in the estimated spoilage are likely, particularly at high demands. Differing spoilage estimates will give large changes in implied demand. Used in reverse, the spill model does not work in practice at high spill. Numerically, the spill model is poor at “detruncation.”

Two methods are used to get around this. Neither are particularly convenient. The simplest is to look at the leg in question at a lower load factor time, and scale the demand up in proportions typical for similar markets suffering less truncation. The second is to collect information from the revenue management system on day-by-day spilled demand and establish the monthly average. To set its levels, revenue management must forecast the unconstrained demand for each fare class for each flight leg. Unfortunately, often these forecasts are discarded. When they are recorded, they are not always very good estimates, because forecasting within revenue management systems also suffers from diminished information when spill is high and past bookings have been capped. Finally, recapture and refill add passengers to observed loads, further complicating the issue. All these complaints aside, estimates from the revenue management system are often the best available.

The overall conclusion on demand estimation is that the spill model is fine for predicting spill when demand is known, but not good at helping with the estimates of demand to begin with.

SUMMARY AND CONCLUSIONS

Spill estimates the demand in excess of capacity. The model assumes a demand distribution and truncates it with a capacity line. The demand distribution is usually Normal. Formulas for a Normal distribution have been used for most applications. Such modeling is widely employed within the airline industry. Revisions to the model recognize an increase in the variance of the demand distribution when demand levels are small, due to random variations adding to the usual cyclic changes in demand. Revisions also adjust for overbooking and revenue management behavior by truncating demand at a capacity somewhat below the physical seat count. The revenue for spilled demand is close to the local discount fare. Spilled demand can be recaptured which diminishes the cost of spill. However, the possibility of refilling with recapture adds value to extra seats. Recapture is small for fleet planning cases and can often be ignored. The great frustration with spill modeling is that for all its effectiveness in estimating spill when the unconstrained demand is known, it can seldom be employed with confidence to unconstrain demand from observed load averages.

Overall, spill modeling has produced practical understandings that have found wide use in the airline industry. Future use may be compromised by rising load factors and a developing trend to use pricing to fill under-utilized capacity. As pricing manipulates demand away from its underlying distribution over flights, the spill model use will require further inventiveness and could become less practical.

REFERENCES

Air Transport Association of America (August 1980), Significance of Airline Passenger Load Factors. Staff Paper.

DeSylva, Enrico (1976?) SPill Model. Boeing Working Paper. [Several references attest to the existence of this paper, but neither the author nor Boeing currently have a copy.]

Schlifer, E. and Y. Vardi (1975) An Airline Overbooking Policy. **Transportation Science**, 9 (2): 101-114

Swan, William M. (1983), Traffic Losses at High Load Factors. **Proceedings of AGIFORS Symposium**, 155-179

