By Wei Wang,
Ravi Kumar, Darius Walczak

DYNAMIC OVERBOOKING: CANCELLATIONS AND NO-SHOWS FOR MAXIMUM REVENUE

Powering Modern Commerce
with Dynamic Pricing Science

EXECUTIVE SUMMARY

Overbooking is a strategy to intentionally accept more booking reservations than the physical capacity of a resource, mainly as an attempt to reduce revenue losses from no-shows or cancellations while taking into account risk of denying service. An overbooking strategy can be considered in any business context where capacity to supply products or services is limited, be it seats in an airline compartment, available rooms of a given type in a hotel, or cars available for rent at a certain car rental location. For ease of presentation, in this white paper we explain all overbooking strategies in the airline context.

Almost all airlines today manage their overbooking in a "static" manner, which is popular in practice mainly due to its simplicity, flexibility, and robustness. With a static overbooking method, once the decision regarding how many seats to overbook is made, the overbooking level is assumed not to change.

In real life, however, the overbooking level should change, for example, if the estimated no-show rate changes or available capacity changes. This typically happens during a re-optimization.

At PROS, we have conducted research on methods that explicitly account for the time dynamics of arrivals and cancellations, as well as seat control, and dynamically make overbooking decisions based on the inventory left and time to departure. We have also demonstrated, through abundant simulations, a consistent revenue benefit that dynamic models have over static models, and we firmly believe that the dynamic models should no longer be overlooked by practitioners.

OVERVIEW

Overbooking is perhaps the oldest revenue management (RM) practice that is still of great importance and widely used in the airline and other hospitality industries today; however, overbooking practiced at most airlines is merely a part of the capacity management process that is fairly static in nature and takes place whenever re-optimization of a flight takes place. Some commonly implemented overbooking methods in practice include but are not limited to the following:

- The "naive" approach ([1, 2]) sets the virtual capacity as the ratio of the physical capacity and the show-up rate in an attempt to ensure that at departure the number of expected show-ups would exactly equal the physical capacity.
- Revenue based overbooking ([1, 2]), which seeks the optimal static overbooking level that maximizes the total expected revenue net of expected denied boarding cost.
- Cost-based overbooking, which determines the optimal static overbooking level based on minimizing the sum of spillage cost (penalty due to denied boarding) and spoilage cost (revenue loss due to flying with empty seats).
- Service-level based overbooking, which determines the optimal overbooking level based on a pre-specified tolerance on number of expected denied boardings.
- Net demand ([3]), which uses gross demand reduced by forecasted cancellations as demand input into the optimization routine so cancellations are not explicitly modeled.

Most static methods are essentially keeping overbooking separate from seat control. More specifically, airlines must first calculate overbooking levels using expected bookings, cancellations, and no-shows, which are then used to establish the number of seats an airline would allow itself to sell (the so-called virtual capacity). After that, some optimization algorithm such as dynamic programming (DP) is used to calculate bid prices (or one of the EMSR heuristics is used to calculate nested booking limits) on this augmented number of seats known as virtual capacity ([4]).

Some dynamic methods allow integration of overbooking into optimization, for example, by including cancellations, no-shows and overbooking in a single-leg DP model as Subramanian et al. ([5]) have done. This discrete-time based model is so comprehensive that there has been renewed research interest recently from both inside ([1, 2]) and outside ([6]) of PROS. With such an integrated DP model, at least in theory, the overbooking problem essentially disappears since the overbooking level is dynamically captured in the bid prices when cancellations, no-shows, and overbooking costs are included in the model.

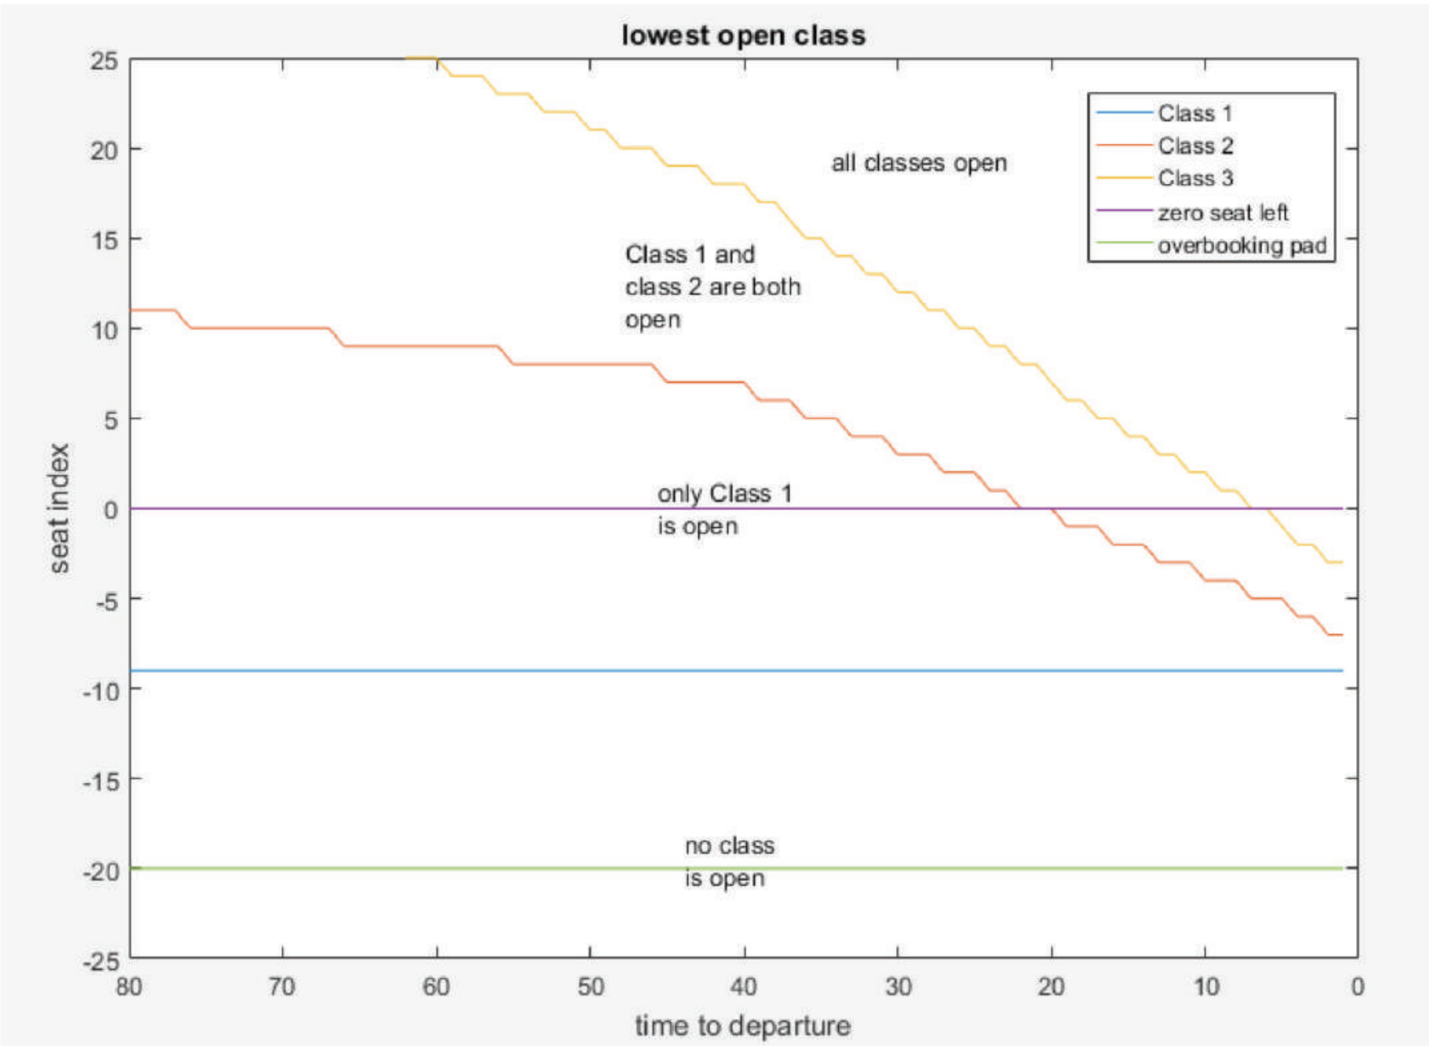

Figure 1: Dynamic Overbooking

As an illustration, Figure 1 illustrates how the overbooking-integrated DP works on a single leg flight with three classes.

In this example, the physical capacity is 25, and the maximum allowed overbooking limit (a.k.a. overbooking pad in the literature) is 20. Note the purple line corresponding to zero seat index together with the green line corresponding to the overbooking pad determine the overbooking range across the entire booking horizon which is equally divided into 80 time slots. More specifically, we can make the following observations from Figure 1.

- Similar to the policy obtained from the traditional DP model, the acceptance/rejection decision can be visually made from the regions defined by the lowest open index trajectories of all the classes. For example, only bookings for the top two classes are accepted in the region defined by the yellow and the orange boundaries.
- Throughout the entire booking horizon, the top class can overbook at most 9 seats (indicated by the blue line).
- When the physical capacity runs out, booking request for Class 2 would never be accepted until it is close to departure (approximately 20 time slots away); in the same situation, Class 3 cannot be overbooked until approximately 5 time slots away from departure.

Note that this example considers no-shows only and the standard DP structural properties still hold. If cancellations are included, some of the structural properties may no longer hold and the lowest open class index lines (e.g. orange line for class 2) may not be monotone.

METHODOLOGY

The traditional DP model has been in the PROS optimization engine for close to two decades. When incorporating cancellations, no-shows and overbooking into the traditional DP model, the following convenient extensions are necessary:

1. The terminal reward (boundary condition for the DP recursion to start) is no longer trivially zero; instead, it now takes denied boarding cost and no-show rate into consideration. We can still make the (currently) typical assumption that each passenger shows up independently with the same probability, thus binomial distribution can be used to calculate the expected overbooking cost as the terminal reward.
2. The fare collected from each booking needs to be properly discounted due to cancellation refund and no-show refund. More specifically, instead of applying the refund at the moment a cancellation or a no-show occurs, one can subtract the expected revenue loss (due to either type of refund) from the fare at the instant of booking, which is known as equivalent charging in the literature [1, 2, 5]. While the cancellation probabilities and no-show probabilities are assumed to be fare class independent in order for the DP model to remain 1-dimensional (and hence efficiently solvable), the cancellation refund and no-show refund are allowed to be class dependent.
3. An additional cancellation term will be included in DP recursion, which is because when a booking gets canceled, the remaining inventory will increase by one, while in the traditional DP, the remaining inventory can only decrease (when a booking requested is accepted) or remain the same (when there is no booking request or the request is rejected). Thus, there are three possible state transitions instead of two (as in the traditional DP).

RESULTS

At PROS, we have conducted extensive simulations comparing dynamic overbooking to its static counterpart under various settings, Figure 2 below includes one such comparison from Wang and Walczak ([3]). In this plot, we used revenue-based overbooking as the static method where the overbooking level stays fixed across all DCPs. We are also aware of the fact that, in practice, overbooking levels are updated at re-optimization times (typically at each DCP), so we introduce the "semi-static" case where the overbooking level is updated at each DCP. We see that clearly dynamic overbooking yields the highest revenue, followed by the semi-static approach, with the static method ranking last.

While, in practice, frequent re-optimization may have an effect on the magnitude of the improvement, we think that there is merit in implementing optimization models with integrated no-shows, cancellations and overbooking provided good forecasts of the required rates can be obtained. Moreover, we have not only seen that the dynamic models consistently lead to higher revenue, but also that, qualitatively, they generally show similar features as the static and semi-static methods. This should be reassuring to practitioners of revenue management.

Refer to Wang and Walczak ([1]) to review more extensive results. In particular, we have observed several phenomena that may not be immediately intuitive at first sight, and which are important for designing airline products with an eye toward ancillary revenues. For example, some types of refund structure may actually help profit or prevent more losses from cancellations and no-shows or, in other words, there are cases where cancellations or no-shows should not be discouraged completely. We have also seen that in addition to the traditional objective of balancing the cost of flying with empty seats against the cost of denied boarding, there can be either positive or negative revenue impact from overbooking even if all booked customers show up.

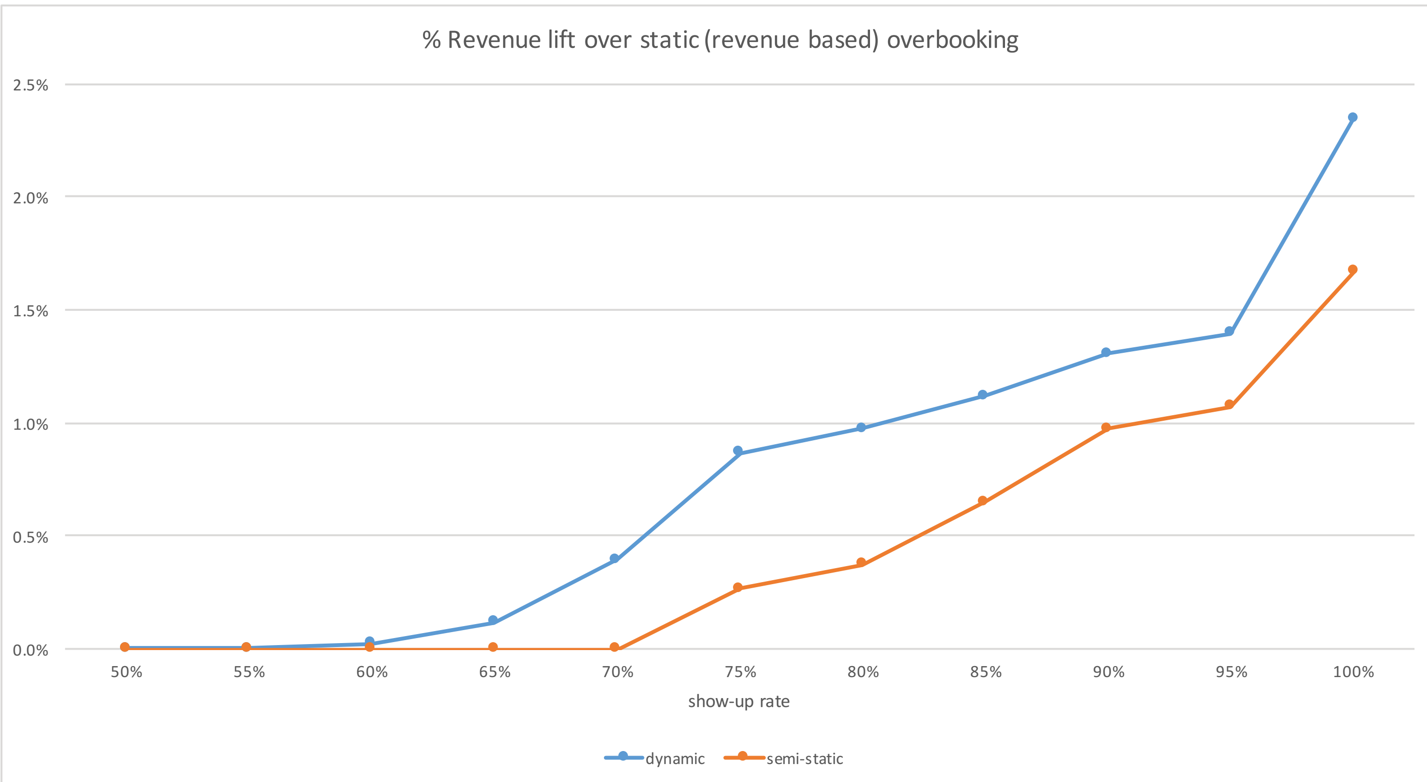

Figure 2: numerical comparisons

REFERENCES

1. Wang, W., Walczak, D. (2016) Overbooking under dynamic and static policies, Journal of Revenue and Pricing Management, 15 (6): 534–553.
2. Wang, W., Walczak, D. (2015) Overbooking under dynamic and static policies, PROS presentation at AGIFORS, Shanghai, China.
3. Li, A., Walczak, D. (2016) Calculating Net Demand for O&D Optimization, PROS Technical Report.
4. Walczak, D., Boyd, E.A. and Cramer, R. (2012) Revenue management. In: C. Barnhart and B. Smith (eds.) Quantitative Problem Solving in the Airline Industry. Boston, MA: Springer, pp 101–162.
5. Subramanian, J., Lautenbacher C. J., and Stidham S. J. (1999) Yield Management with Overbooking, Cancellations and No Shows, Transportation Science, 33: 147-167.
6. Fiig, T. (2016) Joint overbooking and seat allocation for fare families, Amadeus presentation at AGIFORS, Frankfurt, Germany.

ABOUT THE AUTHORS

Wei Wang

is a Senior Scientist at PROS, Inc., where he is one of the key researchers on the company's latest generation of RM platform with two pending patents. He has also been instrumental in the development of the company's suite of visualization tools which helps customers explore and better understand the science behind PROS products. He received his PhD in Operations Research from the University of Washington with his emphasis on stochastic global optimization.

Ravi Kumar

is a Scientist at PROS Inc., in Houston, Texas. His work involves researching Revenue Management and Pricing solutions for the airline industry. He has a Ph.D. in Operations Research from Cornell University where his research focused on the applications of stochastic Dynamic Programming in energy efficient operations of computing and telecommunication systems. Ravi also holds a Master's degree in Mechanical Engineering from the State University of New York at Buffalo where his research involved estimation and control of dynamical systems.

Darius Walczak

is a Principal Research Scientist at PROS, Inc. in Houston, Texas. Since joining PROS in 2001 he has been involved in design and deployment of numerous revenue management (RM) and pricing optimization applications across variety of industries. Darius holds a Ph. D. degree in Business Administration (Operations & Logistics) from Sauder Business School, University of British Columbia in Vancouver, B.C., along with Master's degrees in Mathematics from UBC and Applied Mathematics/Engineering from Wroclaw University of Technology.

About PROS

PROS Holdings, Inc. (NYSE: PRO) is a revenue and profit realization company that helps B2B and B2C customers realize their potential through the blend of simplicity and data science. PROS offers cloud solutions to help accelerate sales, formulate winning pricing strategies and align product, demand and availability. PROS revenue and profit realization solutions are designed to allow customers to experience meaningful revenue growth, sustained profitability and modernized business processes.

To learn more, visit pros.com.

PROS.com

030817