Constrained Demand Forecast

submitted by Air Canada

Context

In AC's current revenue management technology, our system optimizes the inventory available for different city pairs at fixed price points to ensure we maximize the revenue obtained from our fixed capacity. One major part of determining those optimal inventory controls is demand forecasting. Our system forecasts the total potential demand for each city pair (henceforth called unconstrained demand) by offsetting the effects of historical capacity constraints and inventory policies.

After performing the optimization, our system estimates the portion of the unconstrained demand (called re-constrained demand) that will materialize under the optimal inventory policy. This demand is later used by other groups within the company to optimize processes dependent on the real number of passengers having boarded the plane (catering, airport staffing, aircraft weight and balance, etc.).

Problem

Our current re-constraining algorithm has multiple limitations, making it inaccurate for flight date forecast calculations (we do not wish to replicate the optimization logic). Indeed the re-constraining algorithm depends upon the unconstrained demand forecast (very granular data), does not properly reflect additional inventory constraints by AC, and does not consider re-capture (i.e., the purchase of a more expensive seat by a passenger already having a seat). Even though constrained forecast is not a variable optimized or tuned by revenue management, it is a useful guideline for our teams in terms of its impact on demand (a flight expected to leave at 80% is more realistic than a flight forecast to leave at 250% by the unconstrained demand algorithm!). Most importantly, inaccuracy of the forecast has a significant impact in terms of resource waste, unrealized cargo revenue potential, and process misalignment.

Desired solution

Revenue management wishes to obtain an accurate and stable constrained load factor estimation across the life of all our future flights by cabin and type of flow. Special emphasis should be given to the close-in window: we would like to achieve an error that is at most 2% 30 days prior to the flight departure. Here is an example of a partial solution.

On the 31st of January, for a departure on the 14th of March, flight 123 between cities AAA and CCC is expected to have a X% load factor, with X1% of local passengers, X2% of US connections, X3% of domestic connections, and X4% of international connections.

Data

Data will be provided by the revenue management department and include: historical and future revenue controls and inventory states for the flights, current and historical forecast of the flights, current and historical availability and bookings, and current and historical fare forecast.