Updating Bid Price Controls for Dynamic Pricing in the Airline Industry

Gwen Williams[1] Nicola Rennie[1]

[1]STOR-i CDT, Lancaster University

1. Background and Motivation

- Within the airline industry, revenue management aims to maximize revenue by systematically limiting the number of low-cost seats sold.
- Seats are divided into fare classes, with fare class n generating greater revenue the smaller the value of n.

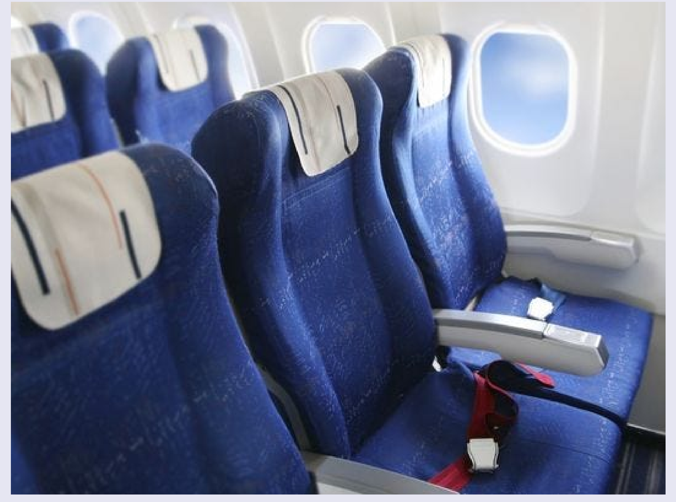

- Bid prices are one method of controlling which seats are sold. They act as a threshold, such that sales are made only if they generate more revenue than the relevant bid price.
- Thus, classes generating revenue lower than the bid price are closed.
- Bid prices are calculated based on **demand forecasts** for each fare class, at different points in the booking period.
- The present research investigated the effectiveness of bid prices when demand **deviates** from forecast, and the potential gains in revenue from **updating** bid prices, according to observed demand.

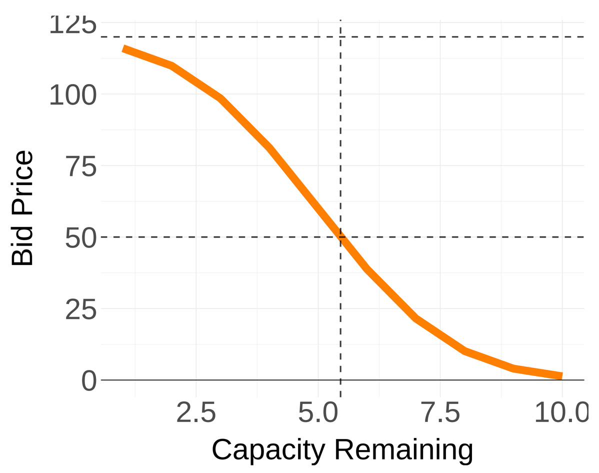

Figure 1: **Static** bid prices are a function of remaining capacity. The vertical line indicates closure of a fare class generating revenue of 50, when remaining capacity is below 6.

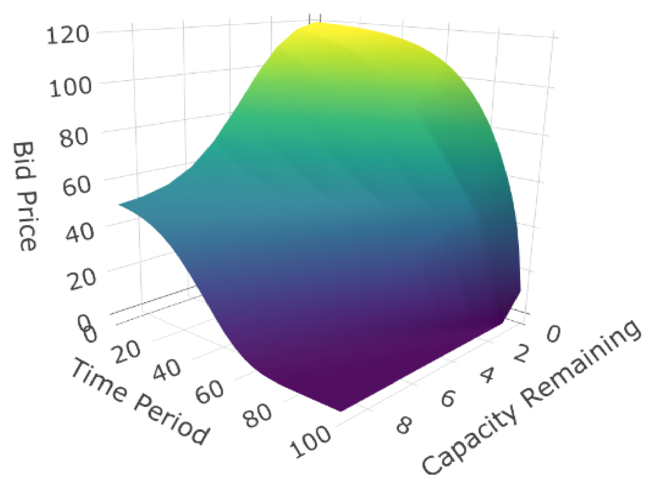

Figure 2: **Dynamic** bid prices are a function of remaining capacity and time-to-flight, where time period 1 indicates the start of the booking horizon.

2. Calculating Bid Prices

This research concerns dynamic bid prices, calculated as the change in value function (optimal expected future revenue) from the sale of the next seat, evaluated at the previous time period. Where,

$$v_t(x) = \sum_{j \in J(x)} p_{jt} \max\{r_j - \Delta_{tj}(x), 0\} + v_{t+1}(x).$$

Here $v_t(x)$ denotes the value function for time t and remaining capacity x, $J(x)$ the set of fare classes, p_{jt} the probability of class j demand during time period t, r_j revenue from class j, and $\Delta_{tj}(x) = v_{t+1}(x) - v_{t+1}(x-1)$.

3. Simulating Demand

- A simple two-class model was used, consisting of two fare classes each corresponding to a group of customers willing to pay a certain amount.
- Bid prices were calculated from an assumed multinomial distribution. Forecasts for the probability of demand arrival during one time period were 0.05 and 0.1 for classes 1 and 2, respectively, for all time periods.

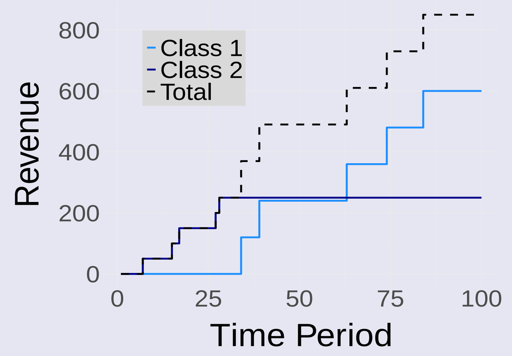

Figure 3: Arrivals were repeatedly simulated for differing demand, and revenue generated from each booking was calculated. The above graph depicts a single simulation.

4. Deviation of Overall Demand

- Ratio of demand between classes remained the same.
- Demand probabilities were homogeneous across the booking horizon.
- Both probabilities increased or decreased from forecast by some factor.

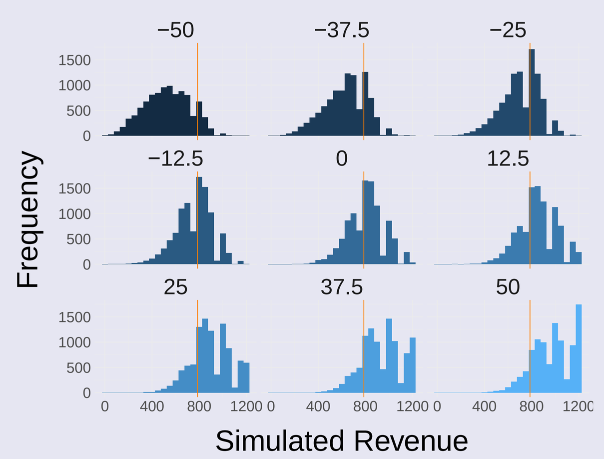

Figure 4: Simulated revenue for demand change of -50% to +50% from forecast, using bid prices calculated from forecast. Orange line indicates average revenue when demand matched forecast.

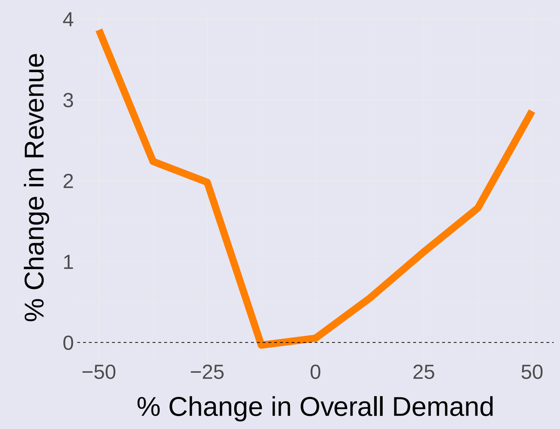

Figure 5: % change in average simulated revenue from **updating** bid prices based on observed demand.

5. Deviation of Demand Over Time

- Ratio of demand between classes remained the same.
- Total expected demand matched forecast, however, demand probability changed mid-way through the booking horizon.

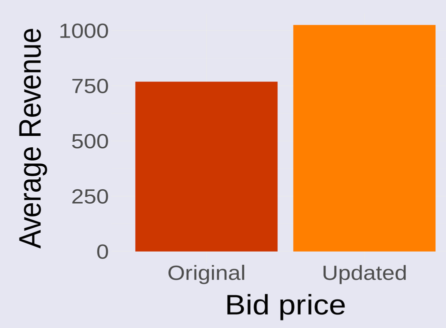

Figure 6: When demand **decreased** over time. Bar chart depicts average simulated revenue when original and updated bid prices were used.

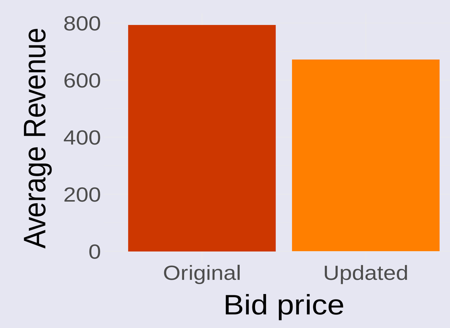

Figure 7: When demand **increased** over time. Bar chart depicts average simulated revenue when original and updated bid prices were used.

6. Computation Times

- The number of calculations required to update dynamic bid prices is $(x \cdot t)^n$, where $x = \textit{total capacity}$, $t = \textit{number of time periods}$ and $n = \textit{network size}$ (a network contains journeys involving a flight leg).
- When x, t, and n are large, very long computation times make it impractical to attempt to update each time demand deviates.

7. Should We Always Update?

Updating does not always have a **significant effect**:

- Updating can only affect revenue by changing which fare classes are open.
- It was found this change often only occurs for a minority of bid prices (see Figure 8).

Updating does not always **increase revenue**:

- Updating bid prices can result in a decrease in revenue (see Figure 7).
- In general, it was found when demand increased over the booking horizon, more revenue was generated by failing to update.

Not all updated bid prices will be **required**:

- Generally, it was found bid prices for low and high remaining capacity around the start and end of the booking horizon, respectively, were rarely ever required (see Figure 9).

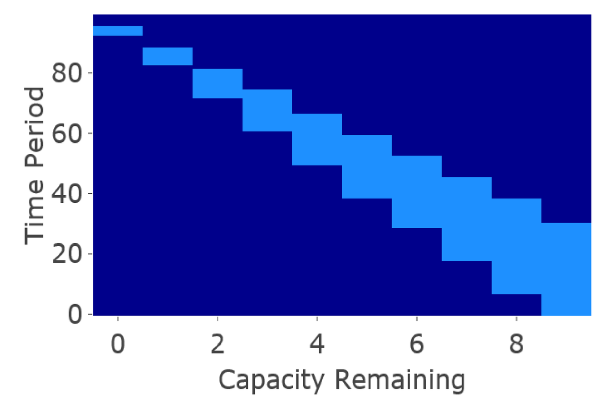

Figure 8: Each coordinate represents a bid price which was updated according to a 50% increase in overall demand. Only those coloured light blue changed which fare classes were open when updated.

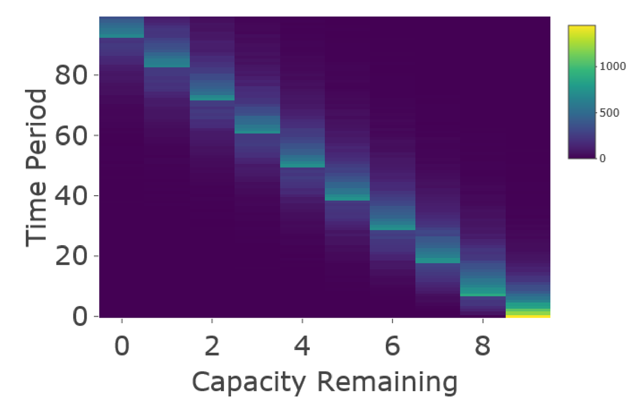

Figure 9: Heat map indicating which bid prices were most frequently referenced over 10,000 simulations, when demand matched forecast.

8. Main Conclusions

1. For substantial **deviations in demand** from forecast, bid prices are **no longer effective** in maximising revenue.
2. Updating bid prices according to observed demand can, in cases, **increase revenue**.
3. However, frequently re-calculating bid prices is **impractical**.
4. Future research should explore possibilities of **selectively updating** bid prices, only in cases where updating increases revenue, and only for bid prices which are likely to be actually required.

http://www.lancaster.ac.uk/stor-i

g.c.williams@lancaster.ac.uk