Adaptive Model Selection Framework: An Application to Airline Pricing

Naman Shukla* [1 2] Arinbjörn Kolbeinsson* [3] Lavanya Marla [1] Kartik Yellepeddi [2]

* Equal contribution [1] University Of Illinois Urbana-Champaign [2] Deepair Solutions [3] Imperial College London

Moving Commercial Decisions from Average Customer Segment to Individual Traveler

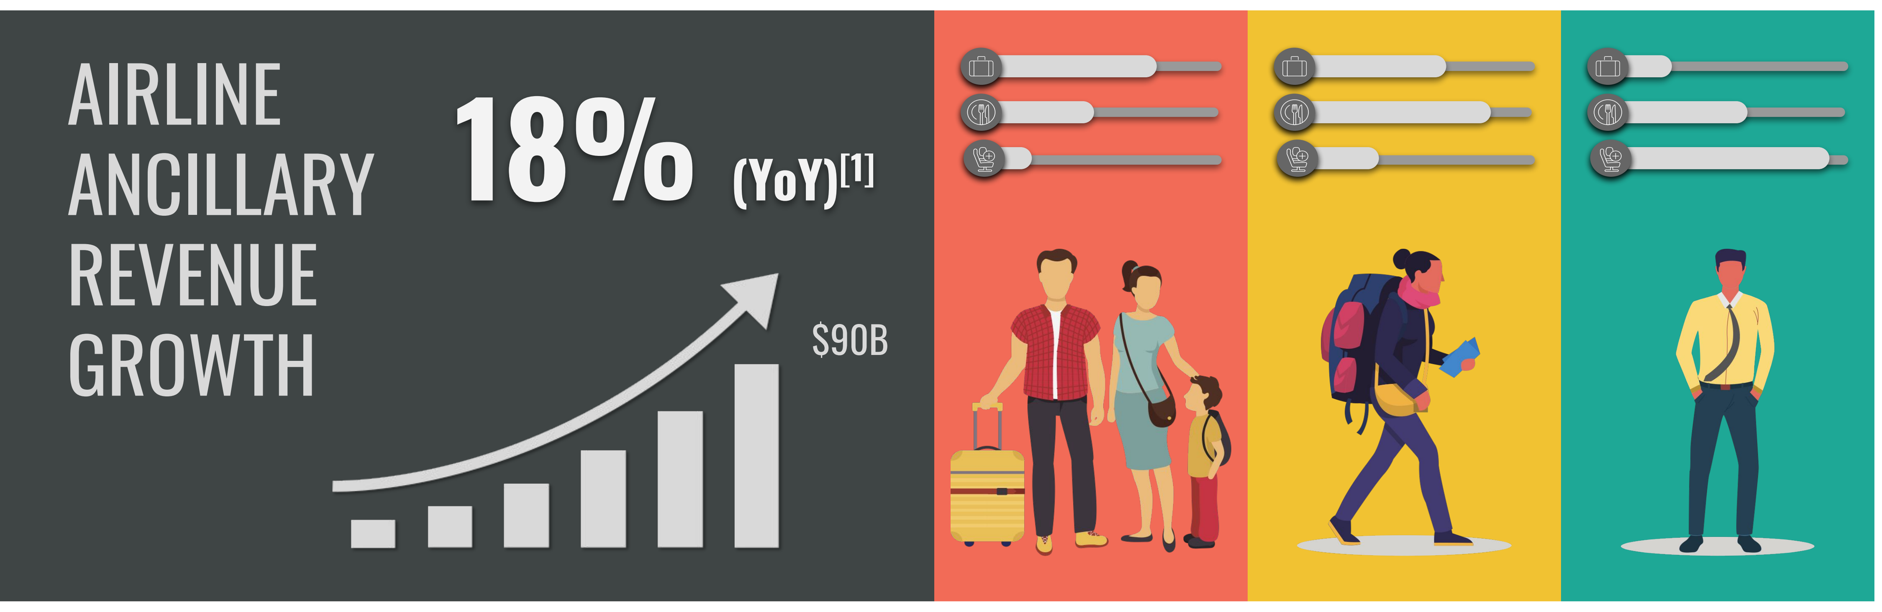

Our Approach: Thompson Sampling (a Multi-Armed Bandit Implementation)

1. Each arm is a different pricing model.
2. The priors are assumed to be beta distributed as shown in the equation and true distributions are assumed to be stationary.
3. During training we update the model selection probabilities as described in Algorithm 1 until convergence.

$$p_{beta}(\theta_a) = \frac{\Gamma(\alpha_a + \beta_a)}{\Gamma(\alpha_a)\Gamma(\beta_a)} \theta_a^{(\alpha_a - 1)} (1 - \theta_a)^{(\beta_a - 1)}$$

Algorithm 1 Thompson sampling for the Bernoulli Bandit

for $t = 1, 2, \ldots$ **do**
  for $a = 1, \ldots, |A|$ **do**
    Sample $\hat{\theta}_a \sim \text{Beta}(\alpha_a, \beta_a)$
  end for
  $x_t \leftarrow \arg\max_a \hat{\theta}_a$
  Apply x_t
  Observe y_t, r_t
  $(\alpha_{x_t}, \beta_{x_t}) \leftarrow (\alpha_{x_t} + r_t, \beta_{x_t} + 1 - r_t)$
end for

A Deployment Framework Designed for Continuous Experimentation[2]

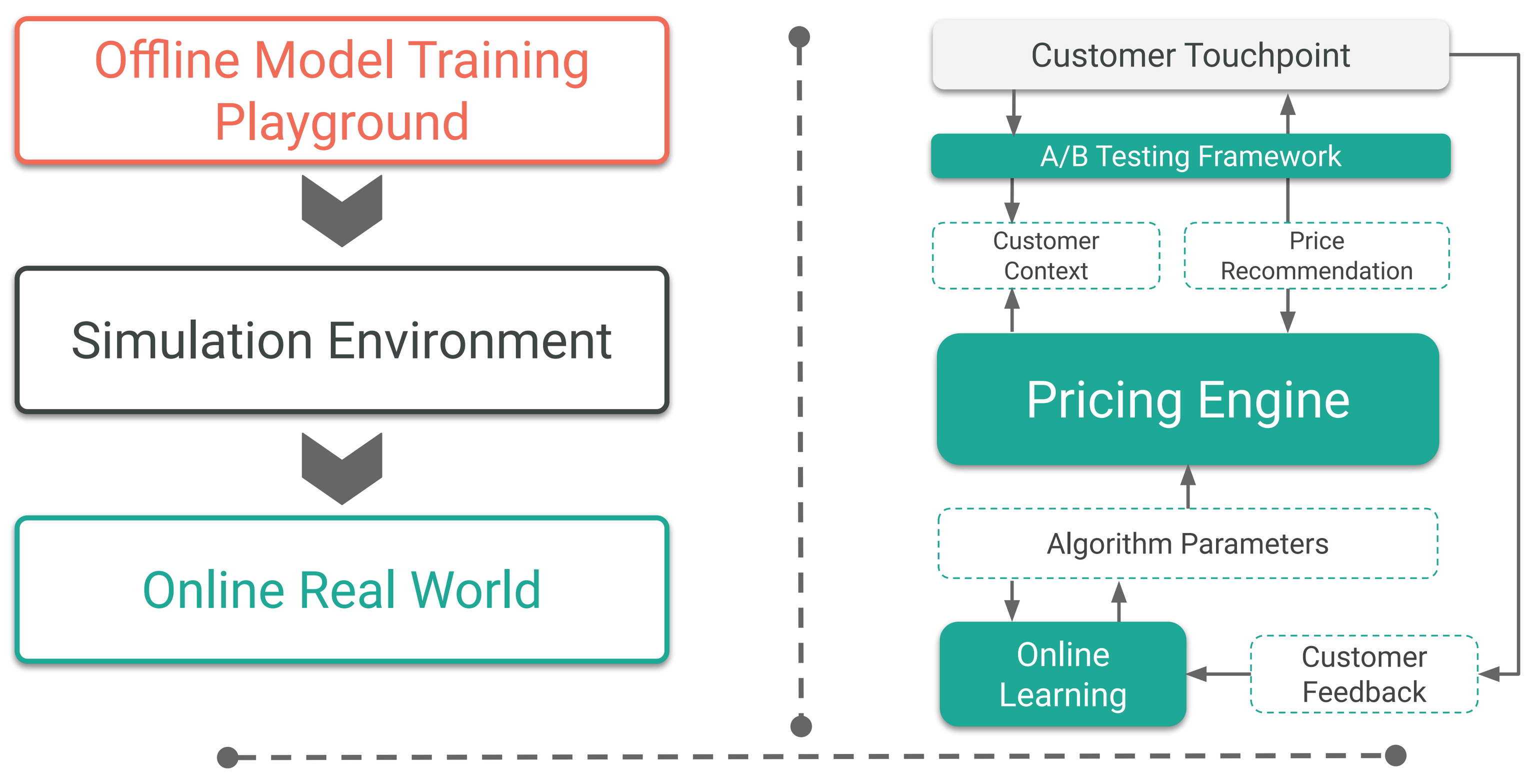

50% Website Traffic

30% Better predictions

17% Higher Conversions

25% Higher Revenue

Results and Discussion

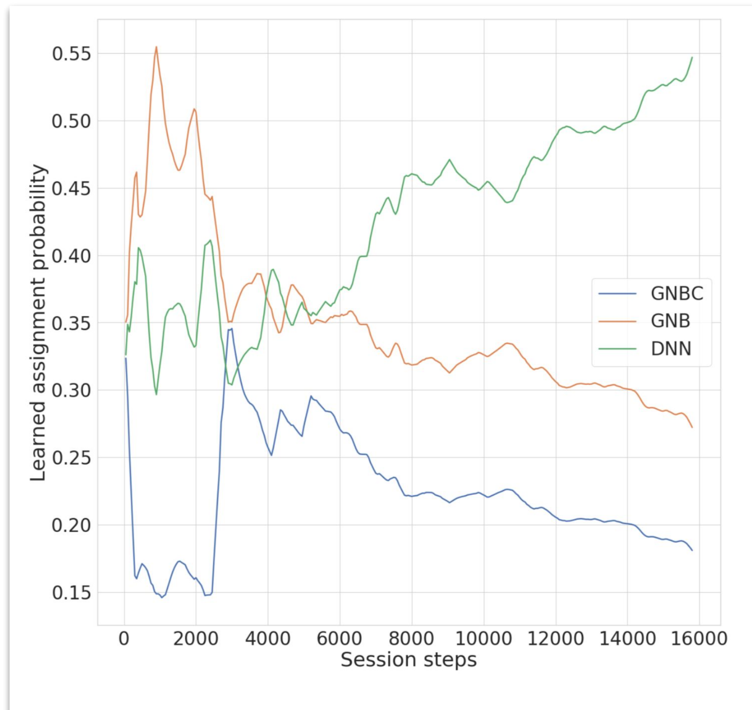

Learned Assignment Probability

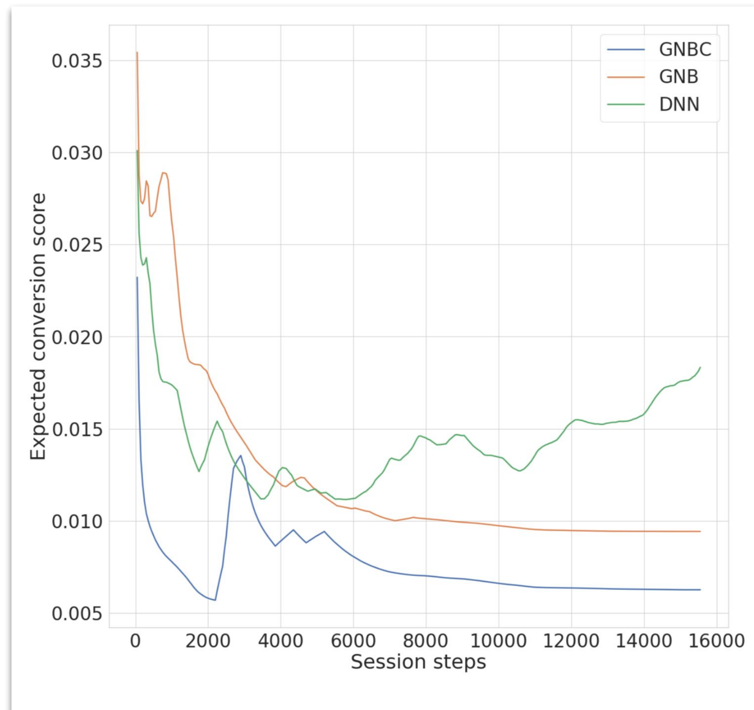

Conversion score of each arm using Thompson sampling

Table 2. Thompson sampling converged values for ancillary pricing using simulated environment. The offer conversion score is the percent of offers generated that were accepted by customers.

Model	Assignment probability	Offer conversion score
GNBC	0.177	0.627%
GNB	0.266	0.943%
DNN	**0.557**	**1.976%**

Table 3. Revenue per offer and Offer conversion scores for different models

Models	Revenue per offer	Conversion score
Only GNBC	0.57	0.61%
Only GNB	0.84	0.93%
Only DNN	1.41	1.71%
Random	0.93	1.00%
MAB	1.33	1.58%

Key Takeaways

1. We develop an approach that uses multi-armed bandit method to actively and adaptively route pricing requests to multiple models to further improve revenue.

2. We test this approach in a rigorously constructed simulation environment to demonstrate that a improved routing scheme that improves business metrics can be achieved.

3. We lay a foundation for future research to use contextual multi-armed bandit methods and perform online testing of this approach.

Our Model Selection Problem and Set-up

Our Pricing Models

1. **GNB:** Gaussian Naive Bayes model for ancillary purchase probability prediction and a pre- calibrated logistic price mapping function for revenue optimization.
2. **GNBC:** Gaussian Naive Bayes with clustered features model for ancillary purchase probability prediction and a pre-calibrated logistic price mapping function for revenue optimization.
3. **DNN**: Deep-Neural Network (DNN) trained using a weighted cross-entropy loss function for ancillary purchase probability estimation. For price optimization, we implement a simple discrete exhaustive search algorithm that finds the optimal price point within the permissible pricing range.

Simulation Environment

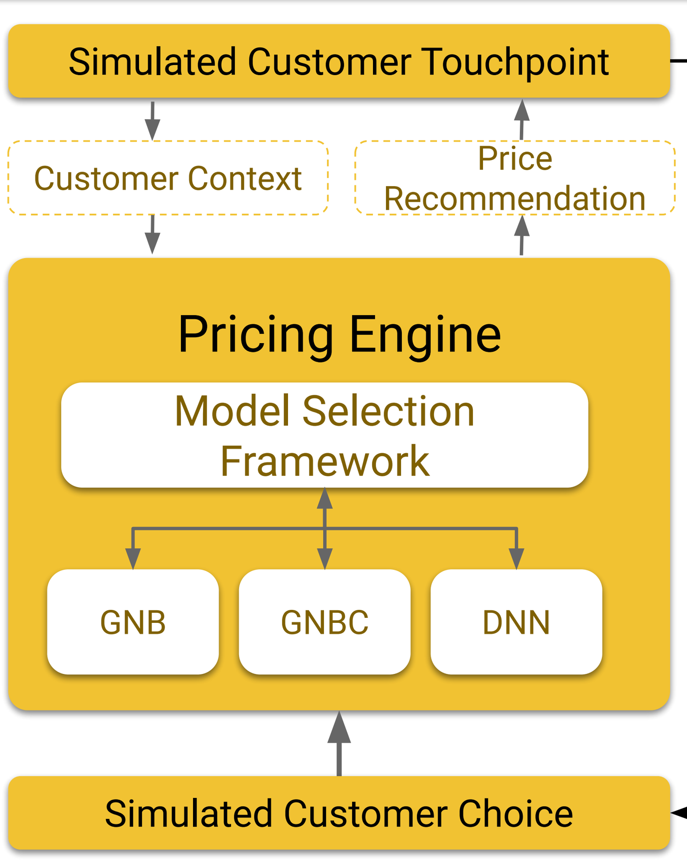

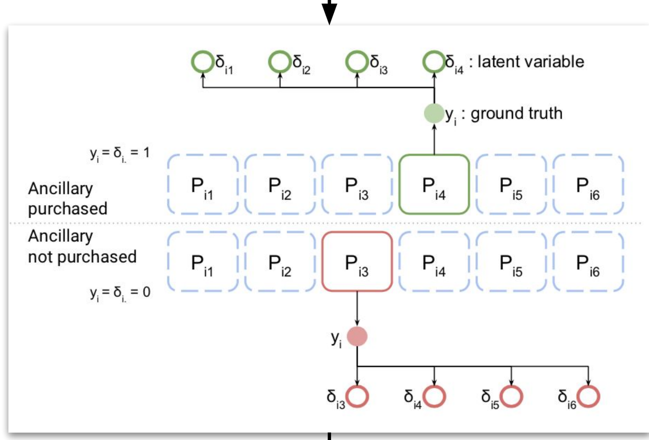

Future Work

Contextual Multi-Armed Bandit: Right now we are tracking the performance of each model at the highest possible level. It is possible that certain models perform better in certain customer contexts. Contextual multi-armed bandit forumation will help us address this issue.

Non-Stationary Formulation: Changing competitive environment and market dynamics could change customer behavior and the performance of various models in these conditions. We plan to explore a non-stationary formulation of the multi-armed bandit, which ignores historical observations made prior to a certain time [3].

Different Objective Function: Currently improvement to the revenue per offer is entirely implicit but not guaranteed. In the future, we plan to adjust the formulation of the multi-armed bandit method to maximize for revenue per offer directly.

References

1. IdeaWorks. Airline ancillary revenue projected to be $92.9 billion worldwide in 2018. https://www.ideaworkscompany.com/wp-content/uploads/2018/11/Press-Release-133-Global-Estimate-2018.pdf, 2018. IdeaWorks Article
2. Shukla, N., Kolbeinsson, A., Otwell, K., Marla, L. and Yellepeddi, K., 2019. Dynamic Pricing for Airline Ancillaries with Customer Context. arXiv preprint arXiv:1902.02236.
3. Cortes, C., DeSalvo, G., Kuznetsov, V., Mohri, M., and Yand, S. Multi-armed bandits with non-stationary rewards. CoRR, abs/1710.10657, 2017.

This Paper

Related Research [2]