# From Average Customer to Individual Traveler: A Field Experiment in Airline Ancillary Pricing

Naman Shukla, Arinbjorn Kolbeinsson
Deepair Solutions, naman,arinbjorn@deepair.io

Lavanya Marla, Akhil Gupta
University of Illinois at Urbana-Champaign, lavanyam,akhilg3@illinois.edu

Kartik Yellepeddi
Deepair Solutions, kartik@deepair.io

**Motivation.** In a machine-machine world of commerce, travel suppliers like airlines will need AI tools that can predict an individual customer's needs during online interaction with the customer, and use those predictions to make commercial decisions that meet customer needs and maximize revenue. Deep Learning (Hinton et al. 2012, Gu et al. 2018, Krizhevsky, Sutskever, and Hinton 2012, Simard et al. 2003) and reinforcement learning (Mnih et al. 2015, 2016, Lillicrap et al. 2015) show promise of being able to find intricate patterns in data by using pattern recognition and machine learning.

**Problem Statement.** In this work, we examine this idea on one product sold by a travel supplier (an airline ancillary). We demonstrate how the supplier can use AI to personalize one commercial decision (pricing) at the point-of-sale, in a manner that is both relevant to the individual customer and profitable to the supplier. Even though this revenue stream is clearly substantial to the airline industry, its pricing strategies are not fully developed due to its recent emergence in the market. Because these products were traditionally not offered separately but bundled along with the flight, airlines have little knowledge of the relationship between the customers' choice of primary product and the ancillary product. Moreover, given the now optional nature of these products, ancillary purchases reflect deep personal preferences of each individual and the context of their trip. Consequently, airlines experience very low conversion rates (less than 5%) for ancillaries. Understanding these personal preferences based on the context of each shopping session is crucial to pricing them effectively. Conventional pricing frameworks are static, and not capable of recommending a price conditioned upon rich session-specific information. We study the question of designing AI models that learn customer preferences, and the question of choosing between multiple available models on an A/B testing framework.

**Approach.** We aim to first estimate the demand distribution function $f_\theta(\boldsymbol{x}, P)$, where $\boldsymbol{x} \mid \boldsymbol{x} \subseteq \mathbb{x}$, is the feature vector and $P$ is the offered price. We used over 30 features that fall under the following categories: (i) Temporal features : Length of stay, seasonality (time of the day, month of

the year, etc), departure time, arrival time, time to departure and time of shop. (ii) Market-specific features : Arrival and destination airport, arrival and destination city, ancillary popularity for the route, etc. (iii) Price comparison scores : Scores based on alternative/same flights across/within the booking class. (iv) Journey specific features : Group size, booking class, fare group, number of stops, etc.

We implemented two kinds of pricing models, (1) Sequential two-stage pricing model and (2) Single stage end-to-end pricing model. A traditional two-stage pricing model comprises of two sequential sub-models: the ancillary purchase probability model followed by the revenue optimization model. These are the models GNBC-LM and DNN-DES shown in Figure 1. The GNBC-LM model uses a Gaussian Naive Bayes (John and Langley 1995) model with clustered features for ancillary purchase probability prediction, followed by a logistic price-mapping function for pricing. The DNN model uses a Deep-Neural Network (DNN) trained using a weighted cross-entropy loss function for ancillary purchase probability estimation, followed by a Discrete Exhaustive Search (DES) over the integer range of possible ancillary prices. In our third model which is a single stage end-to-end custom pricing model, the demand as a function of the price and the price recommendation are simultaneously solved for, using a customized loss function to achieve the recommended price. The customized loss takes into account a regret of pricing low, conditional on the ancillary being purchased; and a penalty for recommending high, conditional on it not being purchased.

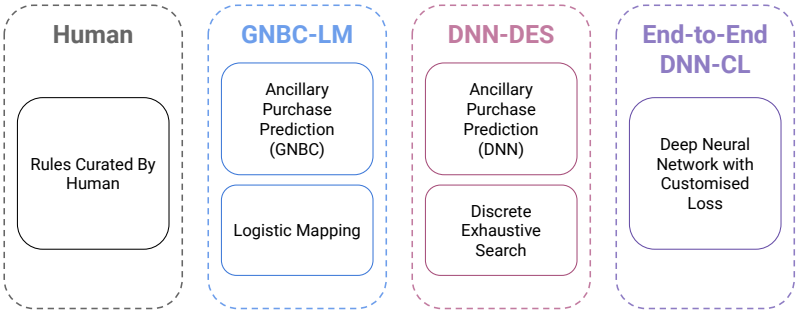



**Figure 1 Schematic of Ancillary Pricing Models**

In order to measure the relative performance of these models they have to be deployed in parallel and connected to the customer touch-point. Often, we find, that no single model outperforms the other deployed models in all customer contexts. Moreover, we are constantly developing newer models that perform well in the offline training playground and show the promise to do better in an simulated or online setting. As a result, we face an exploration-exploitation dilemma – do we (a) exploit the one single model that does best on an average, or (b) explore other models that do better in a different context or (c) utilize offline metrics to choose between models? Therefore,

we design a model selection framework (Figure 2) that addresses this issue by intelligently routing the incoming purchase traffic to various models based on their performance. We model this as a multi-armed bandit problem.

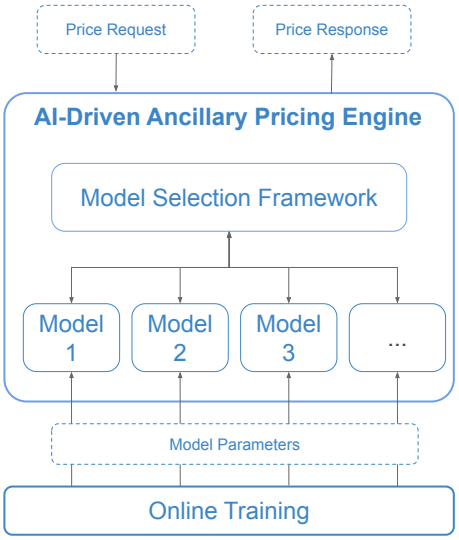



**Figure 2** **Model selection framework within a AI-drive ancillary pricing engine**

In a Bernoulli bandit problem, there are a total of $|A|$ valid actions, where an action $a_t \in A$ at time $t$ produces a reward $r_t \in \{0,1\}$ of one with probability $\theta_a$ and zero with probability $1-\theta_a$. The mean reward $\theta = (\theta_1, \ldots, \theta_{|A|})$ is assumed to be unknown, but is stationary over time. The agent begins with an independent prior belief over each $\theta_a, a = 1, \ldots, |A|$. As observations are gathered, the distribution is updated according to Bayes' rule. The priors are assumed to be beta-distributed with parameters $\alpha_a \in \{\alpha_1, \ldots, \alpha_{|A|}\}$ and $\beta_a \in \{\beta_1, \ldots, \beta_{|A|}\}$, and are updated at each iteration in the Thompson algorithm (omitted for brevity).

**Results** We find a 17% improvement in conversion using AI driven pricing engine as compared to human-curated pricing ($P$-value of 0.0065). Moreover, AI driven pricing resulted in 25% incremental revenue per offer as compared to human ($P$-value of 0.0003). This implies that our model recommends lower prices to targeted customers such that the revenue per offer from our model can still outperform the human-curated pricing system.

Our simulation results using Thompson sampling as a meta learning framework are shown in Table 1. The expected offer conversion scores for each model, where each model is considered as an arm of the multi-armed bandit, converge over the iterations of the Thompson sampling algorithm; and the assignment probability to the DNN-DES model dominates after about 5000 live sessions.

**Conclusions** We successfully demonstrate that ancillaries can be dynamically priced without using any user specific information that violates customer privacy. We compared different AI-based dynamic pricing models and show they outperform human-curated pricing systems by generating

**Table 1 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-LM | 0.177 | 0.627% |
| GNB-LM | 0.266 | 0.943% |
| **DNN-DES** | **0.557** | **1.976%** |

17% higher conversion and 25% more revenue per offer compared to airline's current (human-based) pricing policy. Also, we successfully demonstrate, for a dynamic pricing problem in an offline setting, that a multi-armed bandit approach can be used to adaptively learn a better routing scheme to generate 43% more revenue per offer and 58% more conversion than a randomly selected model, when dealing with an ensemble of models. Our work demonstrates the promise of improved business value through precise, accurate, and continuously updated models for customer demand and associated price sensitivity for ancillaries.

## References

Gu J, Wang Z, Kuen J, Ma L, Shahroudy A, Shuai B, Liu T, Wang X, Wang G, Cai J, et al., 2018 *Recent advances in convolutional neural networks. Pattern Recognition* 77:354–377.

Hinton G, Deng L, Yu D, Dahl G, Mohamed Ar, Jaitly N, Senior A, Vanhoucke V, Nguyen P, Kingsbury B, et al., 2012 *Deep neural networks for acoustic modeling in speech recognition. IEEE Signal Processing Magazine* 29.

John GH, Langley P, 1995 *Estimating continuous distributions in bayesian classifiers. Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*, 338–345 (Morgan Kaufmann Publishers Inc.).

Krizhevsky A, Sutskever I, Hinton GE, 2012 *Imagenet classification with deep convolutional neural networks. Advances in neural information processing systems*, 1097–1105.

Lillicrap TP, Hunt JJ, Pritzel A, Heess N, Erez T, Tassa Y, Silver D, Wierstra D, 2015 *Continuous control with deep reinforcement learning. arXiv preprint arXiv:1509.02971* .

Mnih V, Badia AP, Mirza M, Graves A, Lillicrap T, Harley T, Silver D, Kavukcuoglu K, 2016 *Asynchronous methods for deep reinforcement learning. International conference on machine learning*, 1928–1937.

Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G, et al., 2015 *Human-level control through deep reinforcement learning. Nature* 518(7540):529.

Simard PY, Steinkraus D, Platt JC, et al., 2003 *Best practices for convolutional neural networks applied to visual document analysis. Icdar*, volume 3.