

Dynamic Pricing for Airline Ancillaries with Customer Context

Naman Shukla
University of Illinois at
Urbana-Champaign
Urbana, IL
namans2@illinois.edu

Arinbjörn Kolbeinsson
Imperial College London
London, UK
ak711@imperial.ac.uk

Ken Otwell
Otwell Research
Huntsville, AL
ken.otwell@gmail.com

Lavanya Marla
University of Illinois at
Urbana-Champaign
Urbana, IL
lavanyam@illinois.edu

Kartik Yellepeddi
Deepair solutions
London, UK
kartik@deepair.io

ABSTRACT

Ancillaries have become a major source of revenue and profitability in the travel industry. Yet, conventional pricing strategies are based on business rules that are poorly optimized and do not respond to changing market conditions.

This paper describes the dynamic pricing model developed by Deepair solutions, an AI technology provider for travel suppliers. We present a pricing model that provides dynamic pricing recommendations specific to each customer interaction and optimizes expected revenue per customer. The unique nature of personalized pricing provides the opportunity to search over the market space to find the optimal price-point of each ancillary for each customer, without violating customer privacy.

In this paper, we present and compare three approaches for dynamic pricing of ancillaries, with increasing levels of sophistication: (1) a two-stage forecasting and optimization model using a logistic mapping function; (2) a two-stage model that uses a deep neural network for forecasting, coupled with a revenue maximization technique using discrete exhaustive search; (3) a single-stage end-to-end deep neural network that recommends the optimal price. We describe the performance of these models based on both offline and online evaluations. We also measure the real-world business impact of these approaches by deploying them in an A/B test on an airline's internet booking website. We show that traditional machine learning techniques outperform human rule-based approaches in an online setting by improving conversion by 36% and revenue per offer by 10%. We also provide results for our offline experiments which show that deep learning algorithms outperform traditional machine learning techniques for this problem. Our end-to-end deep learning model is currently being deployed by the airline in their booking system.

CCS CONCEPTS

• **Computing methodologies** → **Supervised learning by classification**; **Cost-sensitive learning**; **Neural networks**; *Modeling methodologies*; • **Applied computing** → **Transportation**.

KEYWORDS

dynamic pricing, airline ancillaries, contextual pricing, deep neural networks, classification

ACM Reference Format:

Naman Shukla, Arinbjörn Kolbeinsson, Ken Otwell, Lavanya Marla, and Kartik Yellepeddi. 2019. Dynamic Pricing for Airline Ancillaries with Customer Context. In *The 25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '19)*, August 4–8, 2019, Anchorage, AK, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3292500.3330746>

1 INTRODUCTION

Ancillaries are optional products or services sold by businesses to complement their primary product [3]. In the airline industry, these services or products can be directly related to a passenger's flight itinerary, such as baggage allowance, leg room, seat upgrades or meals, or may be related to the passenger's overall travel plan, for example, hotel rooms, rental cars, or destination activities. The estimated ancillary revenue collected by major US air carriers was more than \$18 billion in 2015, and \$59 billion for airlines around the world in the same year [12].

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 as ancillaries, 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 are a result of 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 and generating revenue. Moreover, various ancillaries compete with each other's "shelf-space" on the website and wallet-share of the customer, so pricing an ancillary in the context of other ancillaries confounds the problem.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '19, August 4–8, 2019, Anchorage, AK, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6201-6/19/08...\$15.00

<https://doi.org/10.1145/3292500.3330746>

Currently, the majority of ancillary products are static price-points, i.e., invariant to customer or itinerary characteristics. Our aim is to develop a price recommendation system specific to ancillary services, to price these products dynamically based on itinerary-specific information.



Figure 1: User live session as state space

Each *booking session* on an airline’s website can be modeled using the state space represented in Figure 1. The first three steps involve the primary product, which is a set of seats on aircraft connecting the origin to the destination, referred to as *right-to-fly*; while ancillary offerings and corresponding customer choice occur from state 4 onwards. Note that in this work we consider only ancillaries offered during the booking session, and not those booked later, such as adding bags after reaching the airport. The goal of this work is to optimize prices dynamically, while estimating willingness to pay; and therefore we will address the latter case in future work.

Conventional pricing frameworks are static, and not capable of recommending a price conditioned upon rich session-specific information. Our pricing suggestions are generated through an A/B testing framework that directs live booking traffic to various deployed models. The integration specifications for inference and data retrieval for training, with respect to the current pipeline at an airline booking engine, are described in Figure 2.

Our contributions are as follows.

- We present a rich, customized, session-specific dynamic price recommendation system for ancillary services that significantly outperforms existing pricing systems in terms of revenue.
- We develop a deep learning model that effectively estimates purchase probability and *simultaneously* prices the ancillary product, by modeling monotonicity properties of customers’ willingness to pay. This model provides improved revenues

and captures customers’ behavior more accurately than sequential models that combine traditional machine learning (or deep learning) models followed by revenue optimization.

- Our model predicts human choice more accurately than our baseline model, resulting in increased conversion for the ancillary product.
- We implement and test our models on real data, both on historical data and by live testing in an airline’s booking system, and demonstrate both offline and online improvement based on live customer usage statistics.

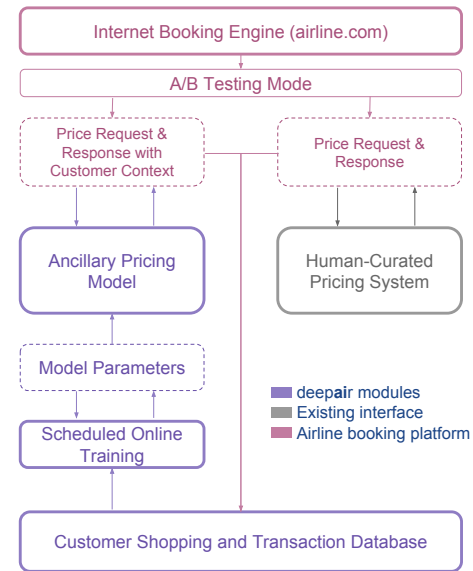


Figure 2: System integration with existing pipeline

1.1 Related Work

Unbundling is the process of separating a product into primary and ancillary products to allow customers more flexibility of purchase, and businesses to increase revenues by matching customer needs more accurately. In the airline industry, this phenomenon has been led by low-cost carriers (LCCs), whose operational and pricing models rely heavily on ancillary fees. In recent years, many legacy airlines have adopted this strategy [10]. Despite an initial response through negative emotions and retaliatory behavior [20], unbundling of services into the basic right-to-fly and additionally priced ancillary services (bags, meals, etc.) has gradually gained acceptance among customers [14]. In fact, revenue from ancillary services in the airline industry have nearly tripled in the past decade, from 3% to 8% of total revenue [19].

Studies on this new phenomena in the airline context are ongoing [1, 6]. Economics literature indicates that the practice of offering add-ons (an equivalent term for ancillaries) can raise equilibrium profits when airlines compete; and can also be used for price discrimination and customer segmentation [8]. Allon et al[1] argue that unbundling and baggage fees are consistent with reduction

of airline operating costs, but may not effectively segment customers. Customer characteristics and the airline's ability to price discriminate are also shown to significantly influence its profits [6]. Bockelie and Belobaba [3] study behavioral models of ancillary product purchase, and specifically comparing the difference in price perceptions of customers who purchase ancillary services sequentially or simultaneously. While multiple behavioral theoretical models [9, 13, 16] based on risk perception, knowledge levels and bounded rationality; and discrete choice models [2] are typically used to model customer choice, there is limited literature that explicitly models the relationship between ancillary services and the primary product (itinerary, or fare class).

Topics in dynamic pricing of homogeneous products have been extensively studied [7]. In fact, dynamic pricing has been a catalyst for innovation in various transport and service industries. Ride-hailing platforms have used surge pricing to match demand and supply, and to avoid the "wild-goose chase" problem [5]. Related to our problem is the work of Ye et al. [21] for Airbnb accommodation pricing. They formulate a custom scheme to optimally price each product using a triple-stage model with booking probability classification, price-suggestion regression and seller-specific logic. Whereas Airbnb considers all their listings as unique and all the customers identical, we consider the inverse problem of identical products and unique customers.

2 PRICING FACTORS

In this section, we discuss the two primary factors we use to determine the optimal price for an ancillary: the demand function and customer attributes.

2.1 Demand Function

An estimation of a demand curve $D(P)$ as a function of price P , can be obtained by evaluating the variation of demand with respect to price. Then, the optimal price can be obtained via maximizing the expected revenue based on the estimated demand curve. The optimal value can only be obtained when the demand function $D(P)$ is an accurate estimate of the actual demand in the market, else the P^* and corresponding revenue will be sub-optimal.

$$P^* = \operatorname{argmax}_P P \times D(P) \quad (1)$$

For airline ancillaries, the demand function D is not just a function of price P offered but also of customer attributes \mathbf{x} . Hence, a better estimation of the demand function is $D(P, \mathbf{x})$ and can be obtained by observing the change in demand conditioned on both price P and customer attributes \mathbf{x} . In our study, we implement and compare two algorithms to estimate the probability of an individual customer purchasing an offered ancillary, which, when aggregated over the population, is a proxy for estimated demand $D(P, \mathbf{x})$. Details of our algorithms are presented in Section 3.

2.2 Customer Attributes

We define a customer's attributes, \mathbf{x} , as the set of factors that influences the probability of that customer purchasing the offered ancillary, *at a given price*. The major attributes that the demand

function is found to be significantly dependent on are time, market, items already in the cart and length of stay.

Time. There are two types of time-related factors that heavily influence the demand function: (1) Days to departure: Price sensitivity captures the relationship between the price and the propensity for purchasing the product. Usually, customers who buy their tickets far in advance are more price sensitive than customers who buy closer to departure. (2) Departure date and time: Like the right-to-fly, ancillary demand has strong time-of-day and seasonal variations. Itineraries starting on certain days and times have higher 'quality' and hence increased demand; due to factors such as higher convenience of time of travel, better connectivity (neither too long nor too short connection time), better availability of alternative connections, special events, and holidays. The quality of service has a strong correlation with the type of passengers it attracts. It is well-known that low quality services tend to be cheaper and attract more price sensitive customers.




Figure 3: Market clusters' probability of ancillary purchase at base price

Markets. Airlines serve a large variety of markets. A market is a tuple comprising of the origin and destination of the trip. Certain markets are served with a larger fraction of non-stop itineraries, while others may be served with larger fractions of itineraries containing connections. Certain markets have a heavier fraction of business trips, while others consist primarily of leisure trips. As shown in Figure 3, there are clusters of markets that have high demand for ancillary services as compared to others. To estimate the demand using (1), we segment these clusters into *sub-markets*. We define *sub-markets* as a mapping from a vector of customer attributes \mathbf{x} to an origin-destination cluster where estimated demand is statistically similar for a given prior ancillary price.

Length of Stay. For those passenger bookings that are round-trips, we define Length of stay (LOS) as the number of days a passenger

plans to stay at the destination. If the passenger does not have a return ticket we consider LOS to be 0. Figure 4 shows estimated kernel density function for LOS for two types of bookings: when an ancillary was purchased (dashed line), and for all bookings (solid line). These estimated graphs are irregular from normalized LOS values of 0.0 to 0.3 (approximately), indicating higher chances of ancillary purchase for LOS durations that are neither too short nor too long. These irregularities indicate that passengers prefer to purchase ancillaries (such as bags), for medium length trips for which they might require additional storage space. Hence, the apparent signature describes the conditional importance of the LOS attribute.




Figure 4: LOS signal from KDE

3 PRICING MODELS

Our pricing model consists of two components: (i) an ancillary purchase probability model that is structured as a binary classification problem, and (ii) a revenue optimization model that, given the probability of purchase, recommends an optimal price that maximizes the airline's expected revenue.

We make the following assumptions:

- **Pricing range:** The recommended ancillary price is allowed to vary only within a legal range defined by business strategy as mentioned in section 6.
- **Monotonicity in willingness to pay:** If a customer is willing to purchase a product for price p , they are willing to purchase the same product at a price $p' < p$. Similarly, if a customer is unwilling to purchase a product at price p , they will be unwilling to purchase at a price $p' > p$.

We implement three different pricing models, of increasing complexity, shown in Figure 5. These are embedded into the framework in Figure 2.

- (1) **Ancillary purchase prediction with logistic mapping (APP-LM):** This model uses a Gaussian Naive Bayes with

clustered features (GNBC) model for ancillary purchase probability prediction and a pre-calibrated logistic price mapping function for revenue optimization.

- (2) **Ancillary purchase prediction with exhaustive search (APP-DES):** This model uses a 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 pricing range.
- (3) **End-to-End DNN with custom loss function (DNN-CL):** This DNN-based model is trained on a customized loss function, and presented in section 3.3. This loss function is designed using the strategic model objective function[21] and explicitly models the dominance properties embedded within the willingness to pay assumption.

In APP-LM and APP-DES, the ancillary purchase probability model and the revenue optimization model are sequential whereas in DNN-CL, they are simultaneously solved to achieve the recommended price.




Figure 5: Schematic of Ancillary Pricing Models

3.1 Ancillary Purchase Probability Model

Our ancillary purchase probability model estimates the demand curve within a sub-market for each offered ancillary, for a given price of the ancillary. This is formulated as a binary classification problem. We aim to estimate the probability distribution function $f_{\theta}(x, P)$, where $x \mid x \subseteq \mathbf{x}$, is the feature vector and P is the offered price. We used over 30 features that fall under the following categories:

- **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.
- **Market-specific features :** Arrival and destination airport, arrival and destination city, ancillary popularity for the route, etc.
- **Price comparison scores :** Scores based on alternative/same flights across/within the booking class.
- **Journey specific features :** Group size, booking class, fare group, number of stops, etc.

As mentioned earlier, the binary classification task is highly challenging because ancillary purchase is highly imbalanced (class ratios of 6 : 100). For APP-LM, We first experiment with many traditional classification algorithms like Gaussian Naive Bayes (GNB), Gaussian Naive Bayes with clustered features (GNBC), Random

Forest (RF), using features chosen based on principal component analysis for these algorithms [15].

For APP-DES, we use a customized deep neural network (DNN) trained on weighted cross entropy loss, as a classifier. While the DNN did not require a lot of feature engineering, we experimented with various hyper-parameters like network architecture, drop-out rates, activation functions, optimization algorithms, and convergence criteria.

3.2 Revenue Optimization

Logistic Price Mapping Function. In our base model APP-LM, once the ancillary purchase probability is predicted, we use a logistic function to recommend a price based on the predicted probability. The intuition behind using logistic mapping is that the ancillary can be priced closer to the maximum of the pricing range when the probability of purchase is high, and lower for low probabilities. Hence, a price mapping is chosen based on (2).

$$p^{rec} = \frac{L}{1 + \exp^{-k(x - x_0)}} \quad (2)$$




Figure 6: Example of logistic mapping from prior probability of purchase to a recommended ancillary price.

According to (2), three parameters can be controlled to map the price desirably.

- Max value, L : this is the full price of the ancillary
- Shape factor, k : the shape or steepness of the curve
- Mid point, x_0 : the mid point of the sigmoid curve

The shape factor k and mid-point x_0 can be fine-tuned to be either aggressive or conservative with pricing. This tuning is illustrated in Figure 6, indicating that at low purchase probabilities, the model compensates by reducing the recommended price.

Discrete Exhaustive Search. Exhaustive search can be efficiently performed over a small set of discrete prices that are within the pricing range. For a given probability of purchase $f_\theta(\mathbf{x}, P)$ and price P , expected revenue is computed using (3).

$$\hat{\mathbb{E}}_P = P \times f_\theta(\mathbf{x}, P) \quad (3)$$




Figure 7: An illustration of a discrete search in the price range

Without assuming that the revenue function is convex, an exhaustive search over all allowed prices can be performed, as represented in figure 7. The convexity assumption is met only when price sensitivity has a small derivative in the region of interest. Thereby, using exhaustive search, the optimal price can be evaluated using equation 4. As discussed in Section 2.1, the optimality of the price p^{rec} is dependent on the accuracy of estimation of the demand.

$$p^{rec} = \underset{p}{\operatorname{argmax}} \hat{\mathbb{E}}_p \quad (4)$$

The performance of a two-stage sequential forecasting and optimization method (like APP-LM and APP-DES), depends on a good demand estimate over the permissible range of prices. This requires sufficient exposure of those prices in the market to learn an accurate price sensitivity curve for each sub-market. Without such data, approximate methods such as custom loss functions can produce more revenue in practice.

3.3 Customized Loss Function for DNN-CL

In this section, we present a customized loss function that 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. The objective function is inspired from the strategic model proposed by Ye et al. [21] and ϵ -insensitive loss used in SVR[17]. We enhance this strategic model using latent variables to incorporate the monotonicity in the willingness to pay assumption in our loss function. Suppose we are given N training samples $\{\mathbf{x}_i, y_i\}_{i=1}^N$, where \mathbf{x}_i is the feature vector and y_i is the ground truth label for the i^{th} session. For purchased ancillaries, y_i equals 1 and 0 otherwise. The recommended price P^{rec} for feature vector \mathbf{x} is denoted by $P^{rec} = \mathbb{F}_\Theta(\mathbf{x}, \mathbb{F})$, where Θ is a set of trainable parameters that can be learned for the mapping function \mathbb{F} , and \mathbb{F} is a set of discrete price points in the pricing range.

The objective of the learning is to minimize the loss \mathcal{L} given as

$$\mathcal{L} = \underset{\theta}{\operatorname{argmin}} \sum_{i=1}^N \sum_{j=1}^{|\mathbb{F}|} (\Phi_{lb} + \Phi_{ub}) \cdot \mathbb{1}_{(\sigma_{ij} > 0)} \quad (5)$$




Figure 8: Latent variable δ mapping from ground truth y

where the lower bound function Φ_{lb} and upper bound function Φ_{ub} are defined as,

$$\Phi_{lb} = \max \left(0, \left(L(P_{ij}, \delta_{ij}) - \mathbb{F}_{\Theta}(\mathbf{x}_i, \mathbb{F}) \right) \right)$$

$$\Phi_{ub} = \max \left(0, \left(\mathbb{F}_{\Theta}(\mathbf{x}_i, \mathbb{F}) - U(P_{ij}, \delta_{ij}) \right) \right)$$

where δ_{ij} , shown in Figure 8, is a latent variable that ensures the monotonicity in the willingness to pay assumption by taking the current ground truth y_i into account. The indicator function $\mathbb{1}_{(\sigma_{ij} > 0)}$ selects loss values corresponding to those δ_{ij} which satisfy the monotonicity condition. Therefore, δ_{ij} is defined as

$$\delta_{ij}(y_i) = \begin{cases} y_i & \text{if } \sigma_{ij} \geq 0 \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

Where, σ is the willingness to pay factor, defined as:

$$\sigma_{ij} = (j - j^*) \cdot (-1)^{y_i} \quad (7)$$

Assuming prices are listed in ascending order, j^* is the index at which P_{ij} equals $\mathbb{F}_{\Theta}(\mathbf{x}_i, \mathbb{F})$. We use L and U for the lower bound and the upper bound of the optimal price range, respectively. The functions $L(P_{ij}, \delta_{ij})$ and $U(P_{ij}, \delta_{ij})$ are defined as follows:

$$L(P_{ij}, \delta_{ij}) = \delta_{ij} \cdot P_{ij} + (1 - \delta_{ij}) \cdot c_1 P_{ij} \quad (8)$$

When the ancillary is purchased, the lower bound L is the purchase price P_{ij} . Otherwise, a lower price of $c_1 P_{ij}$ is set to be the lower bound, where $c_1 \in (0, 1)$.

$$U(P_{ij}, \delta_{ij}) = (1 - \delta_{ij}) \cdot P_{ij} + \delta_{ij} \cdot c_2 P_{ij} \quad (9)$$

The upper bound U is P_{ij} when the ancillary is not purchased, whereas if the ancillary is purchased, a price of $c_2 P_{ij}$ ($c_2 > 1$) is set as the upper bound.

Table 1 illustrates the lower and upper bound loss values for recommended price with respect to discrete price points. For $P_{ij} < \mathbb{F}_{\Theta}(\mathbf{x}_i, \mathbb{F})$, the upper bound loss increases linearly. For upper bound loss to be non-zero, $c_2 > \frac{\mathbb{F}_{\Theta}(\mathbf{x}_i, \mathbb{F})}{P_{ij}}$. Similarly, for non-zero loss, the bounds on c_1 are set to $\frac{\mathbb{F}_{\Theta}(\mathbf{x}_i, \mathbb{F})}{P_{ij}} < c_1 < 1$. For $c_1 = c_2 = 1$, the lower bound and upper bound are equal and hence the optimal

Table 1: Lower bound and Upper bound loss values

| Prices | $\Phi_{lb} \cdot \mathbb{1}_{(\sigma_{ij} > 0)}$ | $\Phi_{ub} \cdot \mathbb{1}_{(\sigma_{ij} > 0)}$ |
|--------------------------------|---|---|
| $P_{ij} < \mathbb{F}_{\Theta}$ | 0 | $\max(0, \mathbb{F}_{\Theta}(\mathbf{x}_i, \mathbb{F}) - c_2 P_{ij})$ |
| $P_{ij} = \mathbb{F}_{\Theta}$ | 0 | 0 |
| $P_{ij} > \mathbb{F}_{\Theta}$ | $\max(0, c_1 P_{ij} - \mathbb{F}_{\Theta}(\mathbf{x}_i, \mathbb{F}))$ | 0 |

price will be the j^{th} price in the price set \mathbb{F} . Therefore, c_1 and c_2 can be chosen to change the gap between the lower and upper bounds.

4 PRICING MODEL EVALUATION

In the absence of optimal price values or the best hindsight strategy, it was important to the airline to define a set of offline and online metrics. Offline metrics are useful for model development, incremental learning, hyper-parameter optimization while online metrics measure business value. Establishing the exact set of offline metrics that correlates with online business metrics is an active area of research (Yi et al. [22]).

4.1 Offline Metrics

In this section, we define the metrics that we use to serve as guides through hyper-parameter tuning and to ensure that nightly update of DNN weights do not overfit the data. We use the *Price Decrease Recall (PDR)* and *Price Decrease Precision (PDP)* scores presented by Ye et al. [21] due to their high correlation with the airline's business metrics. PDR measures how likely our recommended prices are lower than the current offered prices for non purchased ancillary and PDP measures the percentage of non purchased ancillaries where the recommended price is lower than current offered prices. Additionally, we use the following metrics:

Area Under the Curve (AUC). Due to presence of high class imbalance as discussed in Section 3.1, we used the AUC of the Receiver Operating Characteristic (ROC) Curve as the offline metric to compare ancillary purchase prediction model performance.

Regret Score (RS). In recent work [21], regret score has been chosen as an offline evaluation criterion because of its proportional relationship to the business metric. RS is defined by equation (10).

$$RS = \text{mean}_{\text{purchases}} \left(\max \left(0, 1 - \frac{P^{rec}}{P} \right) \right) \quad (10)$$

Intuitively, RS measures on an average how close our recommended price P^{rec} was to the true purchase price P . For the example set of sessions in Table 2, sessions 1, 2 and 5 have 0.20, 0.20 and 0.75 regret respectively. Because sessions 3 and 4 have recommended price higher than purchased, regret is 0. Therefore, RS values for this sample of sessions is 0.095.

Price Decrease F1 (PDF1). This score is inspired by the F1 score used to evaluate the precision and recall trade-off. PDF1 therefore measures the trade-off between PDR and PDP according to (11).

$$PDF1 = \frac{2 \cdot PDR \cdot PDP}{PDR + PDP} \quad (11)$$

Table 2: Example prices for purchased sessions

| Session # | Purchase Price | Recommended Price |
|-----------|----------------|-------------------|
| 1 | 10 | 8 |
| 2 | 15 | 12 |
| 3 | 10 | 15 |
| 4 | 25 | 35 |
| 5 | 40 | 37 |

4.2 Online Metrics

Online metrics represent real-world business metrics that indicate if a model is driving business value. We use two key metrics to measure the performance of a model in the real world.

Conversion Score. One of the primary business metric is the conversion ratio, i.e., the percentage of offers that are being purchased and converted into orders.

$$\text{Conversion Score} = \frac{\text{Number of purchases}}{\text{Total number of sessions}} \quad (12)$$


Revenue per session. The revenue per session (RPS) metric is one of the most essential business metrics to quantify actual performance.

5 TRAINING

All of our deep neural network models (in APP-DES and DNN-CL) are trained on NVIDIA Tesla K80 GPU. We used stochastic gradient descent (SGD)[4] with a decaying learning rate to optimize the loss function. Mini-Batches and drop-out units [11, 18] are used to regularize model training. For a discrete exhaustive search over the prices (see Section 3.2), we use a mini-batch of the allowed price inputs to enable a single call to the GPU which minimizes data transfer and model setup cost for each price selection event. Hyperparameters like c_1 and c_2 are tuned using the bounds for non-zero loss and the upper-lower bound gap (see Section 3.3) over the median price point in set \mathcal{P} . We also use the scheduled mini-batch training approach for online model training.

6 EXPERIMENTS

During the first phase of online experimentation, our airline partner’s business strategy is to test our dynamic pricing approaches in a restricted manner, for a single ancillary, by recommending prices equal to or less than the current human-offered price for that ancillary. Although this strategy reduces the search space considerably, the business motivation behind it is to reduce the overall friction in the traveler’s journey by providing them an incentive to pre-purchase ancillaries online. Therefore, the aim of our online experiment is to offer discounts in an intelligent way, so that we improve the conversion rate of ancillaries without dropping the revenue per offer. This also has a potential negative impact of increasing conversion score without improving revenue per session. Therefore, as mentioned in Section 4, using the right set of metrics for evaluation is crucial.

**Figure 9: ROC curve for offline trained model**

6.1 Offline Experiments

We perform extensive offline experiments to evaluate the performance of the models before deployment. These offline experiments consist of two parts : (i) evaluating classifiers’ performance of ancillary purchase probability, and (ii) pricing effectiveness of a two-stage forecasting and optimization model versus a simultaneous, end-to-end pricing model.

6.1.1 Classifier Performance. It is critical to evaluate the performance of the ancillary purchase probability (APP) classifier since we are using it to estimate the demand function. Due to the high class imbalance present in the data, we used the AUC score to evaluate the performance. We started with Gaussian Naive Bayes with Clustering (GNBC) as our baseline to match the state of the art in the airline industry. We then ran the experiments with the Gaussian Naive Bayes (GNB), Random Forest (RF) and DNN classifiers, all of which performed better than the GNBC baseline. The AUC scores of 0.5716, 0.6273, 0.6633 and 0.7664 for GNBC, GNB, RF and DNN respectively, show that the DNN achieves a 33% improvement in the AUC score compared to our baseline. This improvement is intuitive because DNNs can capture more complex relationships between the input features to predict highly imbalanced classes.

To evaluate the learn-ability and robustness of the classifiers, three datasets with varying amount of data are used. Datasets A, B and C have 41,000, 50,000 and 72,000 sessions respectively. Results from our experiments are in Table 3. DNN shows most dominant signs of learn-ability with increasing dataset size. The best performance of these classifiers on the validation set is also presented as an ROC curve in Figure 9.

6.1.2 Pricing effectiveness of a sequential two-stage model versus a simultaneous end-to-end model. Although the DNN performs well for Ancillary Probability Prediction, it was important to also measure the effectiveness of the final price recommendations from each

Table 3: AUC score of models on datasets.

| Dataset | GNBC | GNB | RF | DNN |
|---------|--------|--------|--------|---------------|
| A | 0.5444 | 0.6013 | 0.6646 | 0.6755 |
| B | 0.5274 | 0.6186 | 0.6771 | 0.6967 |
| C | 0.5716 | 0.6273 | 0.6633 | 0.7664 |

pricing model. We used the offline metrics defined in Section 4.1 to perform the comparison between the two-stage sequential forecasting and optimization models (APP-LM and APP-DES), and the simultaneous end-to-end pricing model (DNN-CL). Given the business requirement to provide discounts on the human-recommended price, we considered Regret Score (RS) and Price Decrease Recall (PDR) as more important than PDP and PDF1 [21]. Our results are summarized in Table 4. The APP-LM (which uses our baseline APP model and is manually tuned through a parameter search), serves as our baseline for pricing effectiveness. The inefficient performance of the APP-DES model on these metrics despite the estimation of a good APP model in the first step suggests that the price-demand relationship (see Figure 7) is not estimated accurately. This shortcoming is overcome by the end-to-end model (DNN-CL), which not only overcomes the effect of this inaccuracy but also outperforms the APP-LM on all four metrics.

Table 4: Comparison of scores for different models in offline experiments.

| Scores | APP-LM | APP-DES | DNN-CL |
|--------|--------|---------|--------|
| RS | 0.0741 | 0.3776 | 0.0726 |
| PDR | 0.6366 | 0.6303 | 0.8294 |
| PDP | 0.9276 | 0.9320 | 0.9230 |
| PDF1 | 0.7550 | 0.7520 | 0.8737 |

Hence, we conclude that the DNN-CL model not only minimizes the regret for not pricing high for purchased ancillaries, but also maximizes the likelihood of the recommended prices being low when ancillaries are not purchased.

6.2 Online Experiments

Our APP-LM model has been deployed in production on our partner airline’s internet booking engine for model validation. The APP-DES and the DNN-CL models are currently under deployment, following their successful performance according to the offline metrics.


According to the airline’s business strategy, we introduced a random discount model in addition to the APP-LM model. This random discount model is allowed to recommend discounted prices based on Gaussian noise. There are two reasons for deploying a random price recommender. First, it establishes a baseline for conversion score improvements from discounted ancillaries. Second, it enables us to explore various prices and calibrate price sensitivity. The deployed models are compared with both human-curated static prices and prices from the random discount model. All three were deployed concurrently in an A/B testing setting for a period of 120

days. The results of this comparison for the most recent 30 days are shown in Table 5.

Table 5: Conversion percentage and revenue generated by our model (APP-LM) compared to human-curated and random prices

| Pricing System | Avg. Revenue per Offer | Conversion Score |
|----------------|------------------------|------------------|
| HUMAN | 1.00 | 10.18% |
| RANDOM | 0.77 | 12.37% |
| APP-LM | 1.10 | 13.92% |

Figure 11 and Table 5 indicate that the random discount model is able to produce higher conversion rates than the human-curated pricing system. This not only demonstrates the existence of price sensitivity among customers, but also allows us to measure it. Additionally, the random discount model is unable to produce higher revenue per offer because it conflates the sub-markets’ demand and makes them indistinguishable, thus losing information.

**Figure 10: Revenue Per Offer by Each Pricing System Over Time**

Because customers’ price sensitivity is observed through the random discount model, a slight increase in the conversion score was expected in our deployed model APP-LM. The results in Table 5, Figure 10 and Figure 11¹ confirm this expectation. We observe a 36% increase in conversion rate, which is a 15% increase compared to the random discount model. More importantly, our model produces 10% more revenue than the human-curated pricing system. This implies that our model recommends lower prices to targeted customers such that the revenue per offer from our model can still outperform (or at least be comparable to) the human-curated pricing system. For revenue per offer and conversion score, we see that our model can indeed capture the market trend in a timely fashion. Furthermore, the clear trend of both higher revenue and higher conversion score with respect to the human-curated system indicate the accuracy of target discount with the customer’s context.

¹The exact dates and revenue figures cannot be included due to proprietary, privacy and sensitivity restrictions




Figure 11: Conversion Score by Each Pricing System Over Time

7 FUTURE WORK

Historically, price sensitivity to ancillaries has not been captured due to static pricing. However, it is critical to capture customers' price sensitivity to price the product correctly and maximize revenues. Currently, all of our proposed models are in early exploratory stage to capture and train on the ground truth responses of customers. While APP-LM has been deployed online, APP-DES and DNN-CL are currently being deployed. Once model validation is performed online, the airline's booking system will switch to the most robust model. For the APP-DES and DNN-CL models, we are specifically interested in further examining the correlation between offline model performance to online business performance, because they outperform APP-LM in offline experimentation. Further, our deployment system will be transitioned from a scheduled mini-batch training (see Section 5), to an event-wise online training. This transition will enable the model to accurately learn temporal dependencies. We also plan to alter the current business strategy (see Section 6) to allow our model to recommend prices higher than current limit, and observe customer responses. Finally, we plan to study the effect of heterogeneous ancillary types being dynamically priced by our deployed models, and the best predicted subset of ancillaries being offered to the customer. Given that various ancillaries compete for wallet-share and shelf-space, it will help expand our understanding of whether such pricing models compete, or collaborate, with each other.

8 CONCLUSION

In this paper, we present a first step in the direction of efficient inference systems from booking data in the airline industry, compared to past works that focus on strategic impacts. We successfully demonstrate that ancillaries can be dynamically priced without using any user specific information that violates customer privacy. We compared three different dynamic pricing models and their associated frameworks. Our results show that the accuracy of estimating the demand and fine-tuning its sensitivity to price, greatly influences the optimality of the recommended price. Our offline experiments show that DNN-CL can perform significantly better than APP-LM, APP-DES, and currently deployed approaches, to maximize revenue. In online experiments, our APP-LM model outperforms human-curated pricing systems. By using reliable evaluation metrics that

correlate well with business impact, we hope to observe further improvement in online metrics through our APP-DES and DNN-CL models that are currently under deployment. 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.

ACKNOWLEDGMENTS

We sincerely acknowledge our airline partners for their continuing support. The academic partners are also thankful to deepair (www.deepair.io) for funding this research.

REFERENCES

- [1] G. Allon, A. Bassamboo, and M. Lariviere. 2011. Would the social planner let bags fly free? (2011).
- [2] M. Ben-Akiva and S. Lerman. 1985. *Discrete choice analysis*. MIT Press, Cambridge, MA.
- [3] Adam Bockelie and Peter Belobaba. 2017. Incorporating ancillary services in airline passenger choice models. *Journal of Revenue and Pricing Management* 16, 6 (2017), 553–568.
- [4] Léon Bottou. 2010. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*. Springer, 177–186.
- [5] Juan Camilo Castillo, Dan Knoepfle, and Glen Weyl. 2017. Surge pricing solves the wild goose chase. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. ACM, 241–242.
- [6] Y. Cui, I. Duenyas, and O. Sahin. 2016. Unbundling of ancillary service: How does price discrimination of main service matter? (2016).
- [7] Arnoud V den Boer. 2015. Dynamic pricing and learning: historical origins, current research, and new directions. *Surveys in operations research and management science* 20, 1 (2015), 1–18.
- [8] G. Ellison. 2005. A model of add-on pricing. *Quarterly Journal of Economics* 120 (2005), 585–637. Issue 2.
- [9] X. Gabaix and D. Laibson. 2006. Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics* 121 (2006), 505–540. Issue 2.
- [10] L. Garrow, S. Hotle, and S. Mumbower. 2012. Assessment of product debundling trends in the US airline industry: customer service and public policy implications. *Transportation Research Part A* 46 (2012), 255–268. Issue 2.
- [11] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*. 249–256.
- [12] IdeaWorks. 2015. Airline ancillary revenue projected to be 59.2 billion dollar worldwide in 2015. <https://www.ideaworkscompany.com/wp-content/uploads/2015/11/Press-Release-103-Global-Estimate.pdf>. (2015). IdeaWorks Article.
- [13] D. Kahneman and A. Tversky. 1979. Prospect theory: an analysis of decision under risk. *Econometrica* 47 (1979), 263–292. Issue 2.
- [14] J. F. O'Connell and D. Warnock-Smith. 2013. An investigation into traveler preferences and acceptance levels of airline ancillary revenues. *Journal of Air Transport Management* 33 (2013), 12–21.
- [15] Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. 1998. Nonlinear component analysis as a kernel eigenvalue problem. *Neural computation* 10, 5 (1998), 1299–1319.
- [16] J.D. Shulman and X. Geng. 2013. *Management Science* 59 (2013), 899–917. Issue 4.
- [17] Alex J Smola and Bernhard Schölkopf. 2004. A tutorial on support vector regression. *Statistics and computing* 14, 3 (2004), 199–222.
- [18] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research* 15, 1 (2014), 1929–1958.
- [19] T. Stalnaker, K. Usman, and A. Taylor. 2016. Airline Economic Analysis. <https://www.oliverwyman.com/content/dam/oliver-wyman/global/en/2016/jan/oliver-wyman-airline-economic-analysis-2015-2016.pdf>. (2016).
- [20] S. Tuzovic, M. C. Simpson, V. G. Kuppelwieser, and J. Finsterwalder. 2014. From free to fee: Acceptability of airline ancillary fees and the effects on customer behavior. *Journal of Retailing and Consumer Services* 21, 2 (2014), 98–107.
- [21] Peng Ye, Julian Qian, Jieying Chen, Chen-hung Wu, Yitong Zhou, Spencer De Mars, Frank Yang, and Li Zhang. 2018. Customized Regression Model for Airbnb Dynamic Pricing. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, 932–940.
- [22] Jeonghee Yi, Ye Chen, Jie Li, Swaraj Sett, and Tak W Yan. 2013. Predictive model performance: Offline and online evaluations. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1294–1302.