Context-Based Pricing for Revenue Optimization with Applications to the Airline Industry

Shivaram Subramanian
IBM Research
subshiva@us.ibm.com

Markus Ettl
IBM Research
msettl@us.ibm.com

Max Biggs
Darden School of Business
University of Virginia
biggsm@darden.virginia.edu

Wei Sun
IBM Research
sunw@us.ibm.com

Youssef Drissi
IBM Research
youssefd@us.ibm.com

Abstract

Most airlines use dynamic pricing to optimize the price of their base economy product by maximizing the expected revenue. However, when it comes to pricing of premium products, airlines often use static price increments that are applied to the best available economy fare based on simple business rules for adjusting the price based on supply. In this paper, we present a suite of machine learning algorithms that take advantage of the rich booking session context available at the time of the booking to make pricing policy recommendations. The challenge is to accurately predict bookings for new combinations of attributes by market and segment (departure time, length of stay, advance purchase, length of haul, ...) using sophisticated machine learning models while keeping the resulting pricing policy interpretable. We present an approach based on a novel path-based Mixed-Integer Programming (MIP) reformulation that can efficiently recover simple yet near-optimal pricing policies. To generate practical pricing policies, the approach accommodates a variety of real-world business requirements into the decision optimization problem. We demonstrate the efficacy of our model with extensive experiments. Finally, we present an airline case study on deriving profitable prescriptive policies for premium cabin tickets based on easily interpretable pricing rules.

1. Introduction

After demand for air travel nearly evaporated at the beginning of the global pandemic in 2020, airlines were forced to cut the supply of flights, dropping routes and cutting frequencies. During that time, most airlines were running at 20 percent or less of their normal flight operations. Now, as we are heading into the holiday season of 2022, travel demand is exploding to levels not seen since before the pandemic and it is outpacing airlines' ability to ramp up capacities to accommodate demand under these changing market conditions.

In such recovery situations, rapid changes are seen in the airline's ticket pricing environment. Historical booking data, which used to be the foundation of revenue management systems while deciding on pricing, have become practically useless. Traditional means of pricing based on time to departure have become irrelevant because most of the bookings seen as the pandemic subsides are being made in a shorter, condensed booking window. Moreover, the new trend of "bleisure" travel where travellers combine leisure activities with professional business obligations on the same trip is expected to further increase in popularity after the pandemic due to changed work habits (e.g., Bove (2022), Schlangenstein (2022)). It will outpace the recovery of business travel, and continue to disrupt the long-established mix of business and leisure bookings that airline revenue managers were accustomed to.

Many airlines are using the recovery phase as an opportunity to be innovative, updating traditional revenue management systems and technology to address the fundamental shifts in the aviation industry that have arisen from the pandemic. Concepts like dynamic and continuous pricing provide the airline with extended flexibility in terms of additional price points on the pricing curve (e.g., Strauss (2021), DeLuca (2021)). A real-time shopping environment with dynamic offers and flexible pricing options helps assessing the level of consumer engagement and enable targeted price responses that are efficient and effective, offering consumers an elevated purchase experience.

In this paper, we describe a suite of machine learning models that provide dynamic pricing recommendations for airline premium cabin products such as premium

URI: https://hdl.handle.net/10125/102801
978-0-9981331-6-4

economy, business class or first class seats. The algorithms take advantage of the rich booking session context available at the time of booking to make pricing policy recommendations while protecting customer privacy. Although most airlines use dynamic pricing to optimize the price of their base economy product by maximizing the expected revenue, premium cabin pricing is often based on static price increments that are applied to the best available economy fare using simple business rules, e.g., Altexsoft (2021).

A practical challenge limiting the adaptation of machine learning models for revenue management is the interpretability of proposed policies. This can prevent practitioners from understanding how prices are derived, leading to a lack of trust in the proposed pricing policy. This is particularly acute for "black-box" models such as gradient boosted trees and neural nets, which have very high accuracy in predicting demand but can result in unpredictable and complex pricing policies which are difficult to verify.

We present an approach based on a novel path-based Mixed-Integer Programming (MIP) formulation, which provides simple, interpretable pricing rules but can efficiently recover near-optimal pricing policies. The models accurately predict bookings for new combinations of booking session features including departure time, length of stay, advance purchase window, or length of haul. We describe a sequential framework where a counterfactual model (teacher) is first trained to predict counterfactual outcomes associated with different actions for every sample in the dataset. Based on the counterfactual estimations, the downstream prescriptive model (student) determines the best set of policies to divide the samples that optimizes the given objective. We use tree-based policies to provide these interpretable policies.

To generate practical pricing policies, the approach is designed to accommodate a variety of real-world business requirements into the decision optimization problem which are paramount for successful enterprise adoption. We demonstrate the efficacy of our model with extensive experiments, and present an airline case study on deriving profitable prescriptive policies for premium cabin tickets based on easily interpretable pricing rules.

2. Related literature

Context-based dynamic pricing is a recent development that's driven by legacy and low cost carriers establishing direct sales channels with travel agencies and the consumer to distribute their flights directly instead of relying entirely on global distribution systems like Amadeus, Sabre or Travelport. Direct distribution allows carriers to control the customers' flow, analyze customer or booking context data, and use their own selling methods. Shukla et al. (2019) develop a deep learning model that recommends personalized prices of airline ancillaries (optional products sold by airlines to complement their primary product, e.g., baggage allowance, seat upgrades or meals) by estimating customer's willingness to pay and optimizing expected revenue per customer. Kosonen (2020) worked with a medium sized network airline to develop context-based pricing policies for ancillary product sales. Mueller et al. (2021) present a dynamic pricing model based on observed online data from consumers with imperfect behavior that efficiently determines equilibrium prices.

Ye et al. (2018) describe a pricing strategy model deployed at Airbnb (the online marketplace for lodging and tourism activities) to set the optimal price for listings of vacation rentals. The method comprises a classification model that predicts the booking probability, followed by a regression model with additional personalization logic that predicts the optimal price for a listing-night combination. The authors also introduce a very useful set of metrics to evaluate the quality of the pricing recommendations, since traditional metrics like precision, recall, F-score or AUC are not meaningful in pricing problems where the ground-truth of optimal price is unknown.

On the methodology side, there has been a recent surge of interest in making machine learning models more interpretable (e.g., Guidotti et al. (2018), Carvalho et al. (2019)). A common approach to infuse interpretability into machine learning models is knowledge distillation where a complex model serves as a teacher and a simpler model learns to mimic the teacher as a student. For instance, in explainable reinforcement learning a black-box teacher model first learns the policies, then a regression tree is trained to approximate these policies (Liu et al. (2018), Puiutta and Veith (2020)). Contrary to such approaches, the teacher in our setting does not produce policy but it provides counterfactual estimations which guide the student in determining the optimal policy by constructing variants of decision trees. The optimal policy can either be constructed greedily Zhou et al. (2018), Biggs et al. (2021) or optimally Kallus (2017), Amram et al. (2020).

As learning an optimal decision tree is NP-complete, popular algorithms such as CART [16] have relied on greedy heuristics to construct trees. More recent advances leverage mixed-integer programming to train globally-optimal trees, e.g., Bertsimas and Dunn (2017), Bertsimas et al. (2019), Aghaei et al. (2019), Aghaei et al. (2020), Zhu et al. (2020). In particular, Bertsimas et al. (2019) and Amram et al. (2020) have applied this MIP approach to learn a prescription policy. Our proposed methodology adapts the column generation-based heuristic approach introduced in Subramanian et al. (2022) to construct prescriptive trees to the context of airline pricing under constrained demands and inter-rule business constraints. Notable differences between our work and the aforementioned MIP approaches are: a) instead of constructing a binary-split tree, we construct a demand-constrained multiway-split tree which make pricing decisions more interpretable; b) instead of an arc-based formulation, we adapt the path-based MIP formulation in Subramanian et al. (2022) to construct the decision tree. Specifically, we adapt and extend the column generation approach to identify (near) optimal pricing policy by jointly optimizing customer segments and the recommended prices for each segment within a single scalable integrated model while satisfying a variety of business and industry requirements such as airline fare filing constraints and practical business rules that eliminate or discourage certain feature combinations. The leaf nodes of the output multi-way split tree represent optimal (counterfactual) pricing actions, and we therefore refer to such a tree as an Optimal Action Tree (OAT) in the remainder of this paper. Next, we briefly describe the predictive counterfactual teacher and prescriptive student machine learning framework that we employ to generate optimal pricing policies.

3. Optimal Action Trees

3.1. Predictive Counterfactual Teacher

With M observational data samples $\{(x_i, \pi_i, y_i)\}_{i=1}^{M}$, where $x_i \in \mathcal{X}^d$ are features, π_i refers to the action chosen from a discrete set Π, and y_i is the uncertain quantity of interest. In the pricing use case analyzed in this work, x_i represents features related to the context of a booking (such as the advance purchase period, departure day of week, one-way or round-trip, class of service, etc.), π_i is the price of a ticket and y_i indicates whether the product was sold or not.

The objective of the predictive teacher model is to provide the counterfactual estimates for each sample across all possible actions, $g_{i,\pi}$, where $\pi \in \Pi$. For the pricing use case, we first train a gradient-boosted tree ensemble model on the observational data using the lightGBM package (Ke et al., 2017) to learn a classifier $f(x, \pi)$ that maps input x_i and π_i to the binary outcome y_i. The function $f(x, \pi)$ is commonly

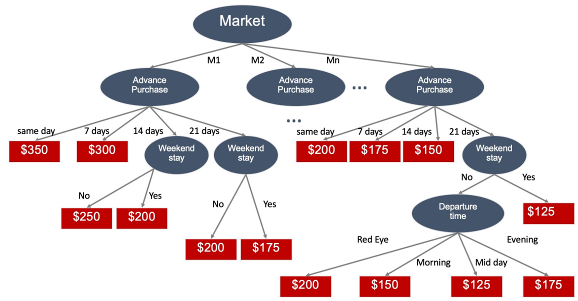

Figure 1: A prescriptive student tree

known as the demand function, which represents the probability of purchase when given price π and booking context features x. The counterfactual outcome we are interested in is the expected revenue for a given price, i.e., $g_{i,\pi} = \pi f(x_i, \pi)$. We want to point out that there are several alternative approaches to training a counterfactual model besides the direct method described here, such as the doubly robust estimator described in Dudik et al. (2011) which requires the propensity scores of historical prices to be estimated.

It is important to note that only a subset of features may be used to define a pricing policy. For example, product availability can be included as a feature in the teacher model but is not valid for defining pricing policy. Customer features such as demographics, geographic indicators, or behavioral features cannot be included within pricing policies due to privacy or fairness concerns. Nevertheless, the teacher model is trained on all available features to provide accurate estimates of the counterfactuals.

3.2. Prescriptive Student Tree

Under the OAT framework, the teacher model estimates the counterfactuals for each sample and the student model constructs the optimal policies in the form of a prescriptive tree by selecting the counterfactuals that maximize a given objective (e.g., expected revenue maximization) subject to constraints. In a prescriptive tree, all samples that are routed to a leaf node are prescribed with the same action. An example of a multiway-split tree is shown in Fig 1. The tree branches on the features “Market” (typically an origin-destination pair), “Advance purchase window”, “Weekend stay”, and “Departure time”. Fig 1 also shows the prescribed price associated with each leaf node.

The high-level idea of constructing such a tree is to first define a feature graph that contains every combination of input features, and subsequently identify a subset of decision rules, n, from the rule space. As

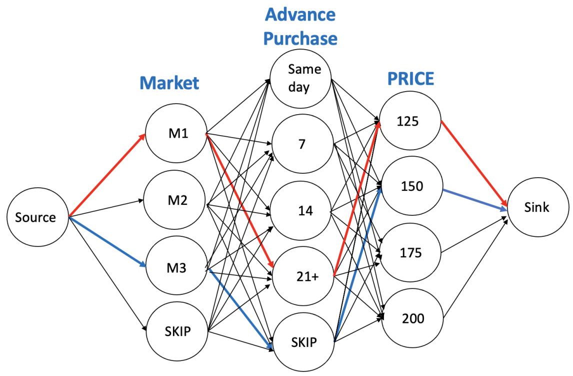

Figure 2: A feature graph with two features and the action (price). Two paths/policies/decision rules are highlighted.

mentioned earlier, only a subset of available features may be used for policy generation. Nevertheless, the rule space may be vast as the number of possible rules is exponential in terms of the feature space. In a feature graph, each feature indicates a level in the graph, represented by multiple nodes corresponding to its distinct feature values. Nodes of one feature are fully connected to nodes in the next level. It also contains a source and sink node. For each feature, with the exception of the action nodes π, there is a dummy node *SKIP*. Fig 2 illustrates an example of a feature graph with two features (i.e., "Market" and "Advance purchase window") and the action nodes ("Price").

For each rule j, define r_j as the corresponding outcome in terms of the counterfactuals, i.e., $r_j = \sum_{i \in S_j} g_{i,p_j}$, where $S_j \in [M]$ is the subset of observations which fall into rule j and $p_j \in \Pi$ is the action prescribed to rule j. The task of selecting a subset of n rules from a feature graph with N feasible paths can be formulated as a mixed-integer programming (MIP) problem:

$$\begin{aligned}
\max \quad & \sum_{j=1}^{N} r_j z_j - \sum_{i=1}^{M} c_i s_i \\
\text{s.t.} \quad & \sum_{j=1}^{N} a_{ij} z_j + s_i = 1, \quad \forall i = 1, \cdots, M \qquad (1) \\
& \sum_{j=1}^{N} z_j \leq n \qquad (2) \\
& z_j \in \{1, 0\}, \quad \forall j = 1, \cdots, N \\
& s_i \geq 0, \quad \forall i = 1, \cdots, M
\end{aligned}$$

where $a_{ij} = 1$ if sample i satisfies the conditions specified in rule j and 0 otherwise. While a sample may fall into several rules, the set partitioning constraint (1) along with the non-negative slack variables s_i that are included with a sufficiently large penalty c_i ensure that each sample is ultimately covered by exactly one rule. The cardinality constraint (2) ensures that at most n rules are active in the optimal solution where n is a user-defined input. The optimal solution of the set partitioning problem corresponds to a multiway-split tree with n rules. It is well-known that the set partitioning problem is NP-hard. Subramanian et al., 2022 propose a computationally efficient algorithm to solve the problem directly via dynamic column generation.

For the pricing use case studied in this paper, additional business constraints must be incorporated to ensure operational feasibility. For instance, one critical requirement for the pricing policy is that the resulting demand do not exceed the available premium seat capacity. Denote $d_{il} = 1$ if sample i is in Market l, 0 otherwise, C_l as the capacity for market l, where $l = 1, \cdots, L$. Recall $f(x_i, p_j)$ is the estimated probability of purchase at price p_j (associated with rule j) by the teacher model. The capacity constraint can be expressed as

$$\sum_{i=1}^{M} \sum_{j=1}^{N} d_{il} f(x_i, p_j) z_j \leq C_l \quad \forall l = 1, \cdots, L.$$

Another constraint common in airline pricing is to ensure that a more restricted airfare must not be more expensive than its unrestricted counterpart, e.g., nonrefundable versus refundable tickets, off-peak departure versus peak time, round-trip ticket versus two one-way fares. Let (j_u, j_r) denote variable indices of pairs of pricing rules whose non-price features are identical except for a single product feature that is unrestricted (u) or restricted (r), and p_U denotes the maximum price allowed. Denote J_u and J_r as the set of rules corresponding to unrestricted and restricted products, respectively. This constraint can be written as the following,

$$\sum_{j_u \in J_u} (p_{j_u} - p_U) z_{j_u} + p_U \geq \sum_{j_r \in J_r} p_{j_r} z_{j_r}, \quad \forall (j_u, j_r).$$

Constraints like the two highlighted are known as the "inter-rule" constraints as they span across several rules. They can be easily handled in the column generation procedure as they influence the subproblem through modified reduced costs to ensure that paths that are more likely to be feasible are added to the MIP. By proactively avoiding rule conflicts, we do not require a separate and suboptimal conflict resolution post-processing step.

Another class of constraints that are important to address in our problem context are "intra-rule' constraints. The scope of these constraints is limited to within a single rule. In many pricing use cases, certain feature combinations are disallowed, e.g., we cannot generate a pricing rule for restricted inventory that is open to advance purchase, and rules with a promotional discount must exclude blackout dates. These constraints are handled within the K-shortest path subproblem as feasibility checks while extending a path in the feature graph. Doing so ensures that the columns that enter the master program are intra-rule feasible.

In the next section, we present modeling enhancements that ensure consistency and stability in our pricing policies, which are critical to user acceptance.

3.3. Stability of Policies

Decision trees are known to be highly unstable, i.e., slight changes in input data may result in drastically different tree structures. Ensemble methods (e.g., AdaBoost, bagging) have been proposed to improve stability, but this usually is at the expense of interpretability. Unstable policies generated from decision trees can lead to challenges in implementing and managing rules and a loss of confidence of stakeholders in the analytics Dwyer and Holte (2007), Li and Belford (2002), and Shannon and Banks (1999). There is no clear consensus on stability measures in decision trees. Existing notions of stability are often defined with respect to generalization error over a loss function, which does not consider the tree structure and content of the resulting policy.

Note that in our Optimal Action Tree based approach, each data refresh (e.g., monthly or weekly) results in a retrained counterfactual estimator whose predicted values may change. These new predicted counterfactuals affect the downstream policies that optimize the expected outcomes which depend on $g_{i\pi}$. As part of this work, we measure the stability of policies and propose a procedure to determine a set of policies that optimizes the underlying objective while maintaining the notion of stability simultaneously.

More specifically, assume there exists a set of existing (active) rules $Z' \in [0,1]^{n'}$ prior to model retraining. Our goal is to determine a new set of pricing rules $Z \in [0,1]^{n}$ which optimizes the underlying objective while remaining as close to Z' as possible. The key idea is to define a metric which measures the similarity between Z and Z', denoted as $\text{Sim}(Z, Z')$, and augment the objective function to encourage similarity between the new and existing policies, i.e.,

$$\max \sum_{j=1}^{N} r_j z_j - \sum_{i=1}^{M} c_i s_i + \lambda \text{Sim}(Z, Z'), \qquad (3)$$

where λ is the regularizer to be determined via cross-validation.

Potential candidates to measure similarity include $L1$ and $L2$ norm. Given that Z is sparse, evaluating the norm is fast. However, such measures ignore the features that make up the policies. For example, consider the following three pricing rules defined by Advance Booking Window (AB), Weekend Stay (WS), and Price:

- Rule 1: AB = 0-6 days, WS = 1, Price = \$350
- Rule 2: AB = 0-6 days, WS = 1, Price = \$300
- Rule 3: AB = 14-20 days, WS = 0, Price= \$175

The L1 and L2 norm on the rules (excluding rule context) will be 0 for all pairs. However, Rule 1 and 2 are more similar in terms of their contexts compared to Rule 3.

Motivated by the soft similarity metric employed in NLP Sidorov et al. (2014), we define a weighted policy similarity matrix $S(Z, Z^{'}) \in R^{n \times n'}$, where

$$S_{ij} = \cos(v_i, v_j) = \frac{\sum_{l=1}^{d} v_{il} v_{jl} w_l}{\sqrt{\sum_{l=1}^{d} v_{il}^2} \sqrt{\sum_{l=1}^{d} v_{ij}^2}}$$

and v_i and v_j of dimension d correspond to the features of active policy i and j from Z and Z' respectively. Feature weights w which are exogenous satisfy $w \geq 0$ and $\sum_l w_l = 1$. Variable weights w capture the different feature importance. Referring back to the airline pricing example, booking context features, flight features, and fare features could be weighed differently. In other words, it may be desirable to impose fewer restrictions on fare prices as opposed to other features with smaller weights w_l for fare features.

By definition, $0 \leq S_{ij} \leq 1$ wherein a score close to zero would indicate no overlap between the policies, and a score of 1 would indicate the policies are identical. The similarity term in the objective function (3) can be replaced by a norm of similarity matrix S. Alternatively, similarity can be incorporated as a constraint in the original formulation.

4. Experiments

We study a problem where we optimize the price of the premium cabin fares to maximize the airline's revenues while holding the main cabin fares fixed. This is due to the existence of a trusted and effective algorithm for setting main cabin prices, but a comparatively rudimentary algorithm for setting premium fares. We use historic transaction data where each entry corresponds to a purchase of a ticket, which is either a standard fare or a premium fare. We use this data to train a teacher model which predicts the probability of an upgrade to a premium cabin seat using gradient boosted trees. This results in an optimization objective where we optimize for marginal revenue gain by multiplying the predicted probability of upgrade to premium, by the difference in price between the premium fare and an opportunity cost, with the premium fare being the decision variable.

The opportunity cost helps to capture inventory effects. If we sell a premium seat, there is one less seat available to sell in the future and also a customer which doesn't purchase a main cabin fare. Our opportunity cost is the maximum of the main cabin fare and an estimate of the expected value of the first-class seat over the remaining selling window which incorporates the limited inventory, which is an input to the algorithm.

While we acknowledge this is a heuristic approach, we also note that the majority of the flights on routes we study are not capacity constrained. This is particularly salient post-pandemic, where 84% of the sales are occurring for seats in the lowest inventory level. Furthermore, for the minority of tickets which face significant inventory pressure, there is a specialized team of analysts which make adjustments as necessary.

The data set used in our experiments is based on six months of booking data from 150 domestic markets served by a major airline. For model training, we extracted a set of over 60 features pertaining to online bookings made on the airline's website. The features fall into the following three categories: booking context (e.g., advance purchase, weekend stay condition, departure day of week, departure time of day, cabin class, seat availability); market features (e.g., hub airport, flight duration, high loyalty tier booking penetration); and pricing features (e.g., standard and premium cabin fares available for purchase at the time of booking).

There is some variability in the historical prices, primarily due to inventory effects that have triggered different prices depending on the scarcity of seats. The inventory level typically isn't revealed to customers and gives us an idea of how customers will respond to different prices. However, we acknowledge this is an ongoing challenge in the airline industry due to the tension between the desire to have consistency of prices over time to meet customer expectations and randomness which facilitates learning demand.

4.1. Impact of feature set on revenue forecast and pricing predictions

A key strength of our teacher-student approach is that the teacher can be trained on a large set of features to estimate counterfactuals and learn accurate customer purchase behavior, whereas the student model that recommends pricing policies might only use a subset of features that the revenue managers select to define the pricing policy. For airlines, the pricing policy is usually based on published fare filing variables such as advance purchase window, one-way or round trip fare, weekend stay, etc.).

In the first experiment, we selected three representative test markets (one medium and two short-haul markets). For these markets, we computed a set of offline metrics to measure the model quality obtained from two runs: one run where the teacher is trained on a reduced set of features used to prescribe the pricing policy, and a second run where the teacher is trained on the entire feature set. In addition to the Mean Absolute Percentage Error (MAPE) of premium cabin revenue predictions, we use the Price Decrease Recall (PDR), Price Increase Recall (PIR), and Booking Regret (BR) scores introduced in Ye et al. (2018) to assess the quality of our pricing recommendations. PDR measures how likely our recommended prices are lower than the airline's historical prices for non-premium cabin bookings; PIR measures how likely our recommended prices are higher than the historical prices for premium cabin bookings; and BR is the median percentage difference of our recommended prices that are lower than the historical prices.

		Market 1 (Short haul)	Market 2 (Medium haul)	Market 3 (Medium haul)
Premium cabin revenue (MAPE)	Reduced	37.1%	29.1%	66.2%
	Full	4.9%	14.4%	15.2%
Price Decrease Recall (PDR)	Reduced	34.6%	5.9%	2.5%
	Full	72.1%	48.1%	48.5%
Price Increase Recall (PIR)	Reduced	80.3%	100.0%	98.0%
	Full	41.6%	63.0%	63.3%
Booking Regret (BR)	Reduced	16.2%	8.9%	16.7%
	Full	6.1%	8.3%	16.1%

Figure 3: Quality of pricing recommendations from runs with reduced and full feature sets.

As illustrated in Figure 3, working with a reduced feature set in the teacher model not only hurts forecast accuracy (over 30% higher prediction error in two of the three test markets), but it also diminishes pricing effectiveness: the high PIR scores combined with low PDR scores indicate that the model trained on the reduced feature set substantially underestimates price elasticity across customer segments, and recommends inflated prices that result in reduced premium cabin conversion rates and lower revenues.

4.2. Business value assessments

The next set of experiments represents a simulated business value assessment based on nearly 500,000 training samples corresponding to booking records spanning eight test markets. The scikit-learn gradient-boosted tree was selected by the automated AI system that we used during experimentation as the best counterfactual estimator. This model predicts the contextual purchase propensity at different price points based on the over 60 features that we extracted for the full feature set. Fifteen price points (within two standard deviations of the historical prices) were considered for recommendations for each test market. These pricing bounds also represent a viable operating range within which the dynamic pricing system can confidently recommend price adjustments.

While the predictive counterfactual teacher model and the student rule optimizer are trained on the same data, the prediction quality as well as expected revenue and pricing metrics are always calculated on a test data set that consists of out-of-sample transactions in a future time period. Also, since interpretability is a critical business requirement, the optimal pricing policy consists of a decision tree of shallow depth (typically a depth of 5 or less), which did not lead to an overfitting problem.

The goal of training the pricing model was to find pricing rules that maximize premium cabin revenues in the test markets, subject to operational constraints such as: minimal rule complexity (pricing rules must be fileable with a central clearinghouse for distribution of fare change information); a limit of no more than ten pricing rules per market; logical restrictions that prevent a pricing policy for a restrictive product (e.g. off-peak departures in a certain market) from being more expensive than an unrestricted counterpart (e.g. any departure time in the same market).

Booking data suggested that demand is close to capacity, which required us to ensure that price sensitive demands do not exceed the available premium seat capacity. This is a simple but critical business requirement that effectively means that the prices across different customer segments in a market have to be jointly optimized and synchronized. We found that an unconstrained (standard) machine learning based approach may do well in prediction, but often fails to generate acceptable pricing policies. In fact, we observed that such an approach can drive up premium cabin demand from early bookers (a segment of customers that is usually highly price sensitive) to more than twice the available capacity, which clearly is not realizable. Moreover, unconstrained price cuts cannibalize premium seat supply for higher value business customers who typically book closer to the departure date.

In direct contrast, our model's ability to satisfy such global constraints enables it to generate prices rules that work in unison to limit the expected demand to within a tolerance of its historical levels. By avoiding steep price cuts for advance purchases, our recommendations also maximize the chances of premium seating availability for high-value customers who often booked close to the day of departure.

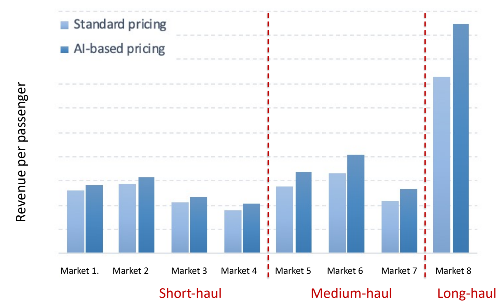

Figure 4: Simulated business value assessment for eight domestic test markets.

As shown in Figure 4, our capacity-constrained pricing policy projects an average 12% gain in premium cabin revenues over the historical levels generated by the airline's legacy pricing practice. The chart compares the simulated premium cabin revenue per passenger (a business metric used by the revenue management team to measure the effectiveness of the pricing policies) before and after optimization for the eight test markets. This gain is achieved by carefully reducing prices for price sensitive segments, along with selectively applying markups for a few late-booking segments associated with a high estimated willingness-to-pay.

A more detailed illustration of the optimized pricing policy is shown in Figure 5. Each bubble in the diagram represents a pricing rule generated by our model. The size of a bubble is proportional to the premium booking revenue generated by customers within the context of the pricing rule.

The fare filing requirements restrict the attributes that can be used within the pricing rules, and this constraint limits the types of customer segments that

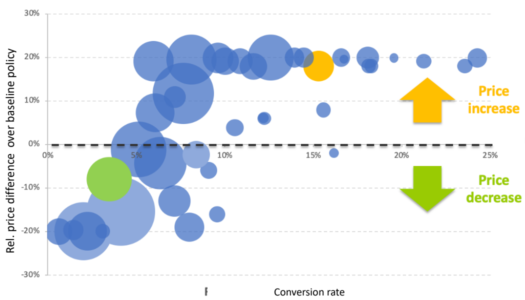

Figure 5: Recommended price changes over the baseline pricing policy.

The results in terms of projected expected revenue gain over the historical (realized) values are shown in Figure 6 for the holiday months. The last column reports the average monthly gain for each of the three scenarios.

Scenario	Month-1	Month-2	Month-3	**Average Gain**
No Retrain and No Re-optimization	-13%	-5%	-30%	**-16%**
Retrain and Stabilized Re-optimization	13%	39%	20%	**24%**
Retrain and Complete Re-optimization	16%	34%	31%	**27%**

Figure 6: Monthly revenue gain during a holiday travel season achieved under different requirements for rule similarity.

can be identified by the column generation procedure to offer differentiated pricing, which in turn can lead to a sub-optimal revenue gain.

4.3. Stability of pricing policies

The final set of experiments analyze the trade-off between rule stability and expected revenue gain. We compare the revenue performance of this stabilized OAT model over a three month rolling window during the US summer holiday travel period. We first identify a good set of weights to calibrate the weighted similarity formula from section 3.2, with the aim of achieving a high similarity score while also ensuring that the projected revenue gain is close to the of optimal value. Next, the predictive teacher model and the OAT pricing rule engine was trained using 3 months of pre holiday season data for a representative medium-haul market to obtain a baseline pricing policy to compare future performance.

We analyze revenue performance for three different stability scenarios. The normalized similarity score for each setting is shown in parentheses:

(a) No Retrain and No Re-Optimization (1.0): The baseline rules are used to score all future transactions.

(b) Retrain and Stabilized Re-Optimization (0.9): The teacher model is retrained with the latest data at the end of each month and the OAT model with the calibrated similarity term in the objective function is reoptimized to generate updated rules and score transactions for the next month.

(c) Retrain and Complete Re-Optimization (0.5): The teacher model is retrained with the latest data at the end of each month and the OAT model (ignoring stability) is reoptimized to generate completely new rules to score transactions for the next month.

We can observe from the first row of Figure 6 that not retraining the teacher model with the latest data negatively impacts performance and results in a revenue loss in each of three months. Hence, it is important to periodically refresh the prediction as well as the pricing policy optimization. Doing so significantly increases the confidence of obtaining positive revenue gains, as depicted in the last two rows. In particular, Month-3, which is a peak travel month in the US, maximally benefits from retraining and complete OAT re-optimization. Independently optimizing rules produces an optimal solution, but the underlying action tree differs significantly from the existing rule set (less than 50% similarity). On the other hand, enforcing similarity does not necessarily diminish revenue performance. By applying the stability framework, the average drop in predicted revenue when compared to a full refresh, is limited to 3% on average. The reduced gain, although undesirable, may be offset by the improvement in user experience and trust due to a more stable model response.

Based on these findings, we recommend that the weights for the similarity formula be optimally tuned as hyper-parameters to limit revenue loss while also ensuring that the changes in pricing policy over time are within an acceptable range to the stakeholders of the pricing system.

5. Conclusions and Future Work

Distilling practically effective prescriptive rules from data is an important emerging area in AI/ML-driven Revenue Management. This work focuses on the airline industry where we employ a novel methodology to identify an Optimal Action Tree (OAT) that unifies the twin tasks of dynamic segmentation and dynamic pricing to achieve superior performance even in a non-stationary post-pandemic environment.

The OAT approach is flexible enough to accommodate customized rule definitions, increase rule stability, and satisfy a wide variety of operational constraints. Our approach can proactively resolve inter-rule conflicts to ensure that the different pricing recommendations across the identified customer segments work in unison to maximize the user KPIs and capacity utilization metrics. We validated the approach in a proof of concept with a global legacy airline, and demonstrate OAT's capabilities through several experiments on a large-scale travel data set. The results reveal superior revenue gains when compared to conventional ML approaches.

In future extensions of this work, we plan to assess the full impact of non-stationarity in the data on customer segmentation and pricing effectiveness. We consider this task to be highly relevant to the travel industry as a whole given the rapid changes in travel patterns and customer purchase behavior observed in recent times.

References

Aghaei, S., Azizi, M. J., & Vayanos, P. (2019). Learning optimal and fair decision trees for non-discriminative decision-making. *Proceedings of the AAAI Conference on Artificial Intelligence*, *33*(01), 1418–1426.

Aghaei, S., Gomez, A., & Vayanos, P. (2020). Learning optimal classification trees: Strong max-flow formulations. *arXiv preprint arXiv:2002.09142*.

Altexsoft. (2021). Dynamic pricing strategy for airlines: Exploring current problems and adopting continuous pricing [Available at: https://www.altexsoft.com/blog/dynamic-pricing-airlines/].

Amram, M., Dunn, J., & Zhuo, Y. D. (2020). Optimal policy trees. *arXiv preprint arXiv:2012.02279*.

Bertsimas, D., & Dunn, J. (2017). Optimal classification trees. *Machine Learning*, *106*(7), 1039–1082.

Bertsimas, D., Dunn, J., & Mundru, N. (2019). Optimal prescriptive trees. *INFORMS Journal on Optimization*, *1*(2), 164–183.

Biggs, M., Sun, W., & Ettl, M. (2021). Model distillation for revenue optimization: Interpretable personalized pricing. *International Conference on Machine Learning*, 946–956.

Bove, T. (2022). Bleisure is the latest work trend, spurred by the pandemic. here is what it means, and why it's being hotly debated [Available at: https://fortune.com/2022/02/12/bleisure-pandemic-work-trend-business/-leisure-travel].

Carvalho, D. V., Pereira, E. M., & Cardoso, J. S. (2019). Machine learning interpretability: A survey on methods and metrics. *Electronics*, *8*(8), 832.

DeLuca, J. V. (2021). Lufthansa group shares the details of its continuous pricing solution.

Dudik, M., Langford, J., & Li, L. (2011). Doubly robust policy evaluation and learning. *Proceedings of the 28th International Conference on International Conference on Machine Learning*, 1097–1104.

Dwyer, K., & Holte, R. (2007). Decision tree instability and active learning. *European conference on machine learning*, 128–139.

Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM computing surveys (CSUR)*, *51*(5), 1–42.

Kallus, N. (2017). Recursive partitioning for personalization using observational data. *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 1789–1798.

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., & Liu, T.-Y. (2017). Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 3146–3154.

Kosonen, E. (2020). *Dynamic Pricing of Airline Ancillaries – Co-creating a Machine Learning Model to Price Ancillaries in the Case Company* (MBA thesis). Metropolia University of Applied Sciences.

Li, R.-H., & Belford, G. G. (2002). Instability of decision tree classification algorithms. *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, 570–575.

Liu, G., Schulte, O., Zhu, W., & Li, Q. (2018). Toward interpretable deep reinforcement learning with linear model u-trees. *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 414–429.

Mueller, D., Nesterov, Y., & Shikhman, V. (2021). Dynamic pricing under nested logit demand. *Pure and Applied Functional Analysis*, *6*(6), 1435–1451.

Puiutta, E., & Veith, E. M. (2020). Explainable reinforcement learning: A survey. *International Cross-Domain Conference for Machine Learning and Knowledge Extraction*, 77–95.

Schlangenstein, M. (2022). Bleisure trips boom as travelers mix work and play [Available at: https://www.washingtonpost.com/business/on-small-business/bleisuretrips-boom-as-travelers-mix-\work-and-play/2022/01/22/bd18443c-7b8b-11ec-9dce-7313579de434_story.html, Accessed on June 26, 2022].

Shannon, W. D., & Banks, D. (1999). Combining classification trees using mle. *Statistics in medicine*, *18*(6), 727–740.

Shukla, N., Kolbeinsson, A., & Yellepeddi, K. (2019). Dynamic pricing for airline ancillaries with customer context. *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.

Sidorov, G., Gelbukh, A., Gómez-Adorno, H., & Pinto, D. (2014). Soft similarity and soft cosine measure: Similarity of features in vector space model. *Computación y Sistemas*, *18*(3), 491–504.

Strauss, M. (2021). Travel technology for dummies: What is continuous pricing? [Available at: https://www.travel-industry-blog.com/corporate-travel/travel-technology-for-dummies-continuous-pricing/].

Subramanian, S., Sun, W., Drissi, Y., & Ettl, M. (2022). Constrained prescriptive trees via column generation. *Proceedings of the AAAI Conference on Artificial Intelligence*.

Ye, P., Qian, J., Chen, J., Wu, C.-h., Zhou, Y., De Mars, S., Yang, F., & Zhang, L. (2018). Customized regression model for airbnb dynamic pricing. *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 932–940. https://doi.org/10.1145/3219819.3219830

Zhou, Z., Athey, S., & Wager, S. (2018). Offline multi-action policy learning: Generalization and optimization. *arXiv preprint arXiv:1810.04778*.

Zhu, H., Murali, P., Phan, D., Nguyen, L., & Kalagnanam, J. (2020). A scalable mip-based method for learning optimal multivariate decision trees. *Advances in Neural Information Processing Systems*, *33*, 1771–1781.