CARMA 2026

8th Int. Conf. on Advanced Research Methods and Analytics
1-3 July 2026, Valencia, Spain
DOI: https://doi.org/10.4995/CARMA2026.2026.21863

AI-Driven Ancillary Pricing: Promise, Assumptions, and Gaps

Feryaal F. Ahmed[1], **John G. Wilson**[1]
[1]Ivey Business School, Western University, Canada

How to cite: Ahmed, F.F.; Wilson, J.G. (2026). AI-Driven Ancillary Pricing: Promise, Assumptions, and Gaps. In: 8th International Conference on Advanced Research Methods and Analytics (CARMA 2026). Valencia, 1-3 July 2026. https://doi.org/10.4995/CARMA2026.2026.21863

Abstract

Artificial intelligence (AI) is increasingly applied to pricing and revenue management, yet its role in ancillary pricing remains fragmented. This paper reviews AI-based approaches in ancillary pricing and organizes the emerging literature by the level at which learning is integrated into pricing decisions: predictive support, predict–then–optimize frameworks, and adaptive policy learning. While supervised models improve demand estimation, only a limited number of studies embed learning directly within pricing policies. We show how AI methods shift modeling assumptions from explicit structural demand specification toward reliance on data, experimentation, and computational infrastructure. Despite methodological progress, the literature remains structurally partial, with limited integration of primary–ancillary dependence, capacity constraints, and adaptive learning. These gaps highlight substantial opportunity for unified AI-enabled ancillary revenue management.

Keywords: *Ancillary Pricing; Revenue Management; Artificial Intelligence; Online Learning; Dynamic Pricing*

1. Introduction

Artificial intelligence (AI) has rapidly transitioned from experimental applications to being embedded in firm-level decision making across industries, enabling data-driven optimization in complex and uncertain environments. Within pricing and revenue management, these advances have prompted renewed interest in learning-based approaches that complement or extend classical optimization frameworks (Yeoman, 2025). While AI adoption in revenue management remains nascent, its potential to support adaptive and personalized pricing is widely recognized (Vinod, 2023).

Among revenue management problems, ancillary pricing presents a particularly demanding setting. Ancillary products are conditional on the purchase of a primary product, creating a hierarchical pricing structure rather than independent pricing decisions. The interaction between primary and ancillary prices under capacity constraints further complicates optimization

Article in Press

This work is licensed under a Creative Commons License CC BY-NC-SA 4.0
Editorial Universitat Politècnica de València

(Wilson & Ahmed, 2024). As additional ancillaries are introduced, the problem scales combinatorially, and customer heterogeneity increases dimensionality (Ødegaard & Wilson, 2016). Analytical solutions with reasonable assumptions have only been for one ancillary item (Ødegaard & Wilson 2016, Wilson and Anderon, 2026). Greater segmentation enhances personalization but intensifies data sparsity and curse-of-dimensionality effects.

These characteristics align closely with the types of problems for which AI methods have been proposed. Supervised learning models can estimate purchase propensities across large feature spaces without explicit functional forms (Bajari et al., 2015), representation learning approaches can capture combinatorial interactions without enumerating all combinations (Peng et al., 2021), and reinforcement learning and bandit frameworks enable adaptive experimentation by updating pricing policies as new information arrives (Bubeck & Cesa-Bianchi, 2012).

Recent reviews synthesize the broader ancillary pricing literature (Mumbower, 2025; Wilson, 2025). However, the integration of AI within ancillary pricing remains limited and dispersed. This review focuses specifically on peer-reviewed studies applying AI or learning-based methods to ancillary pricing environments, identified through Google Scholar and Web of Science. Because this stream of research is recent and relatively small, it permits a focused conceptual synthesis. We organize the literature according to the level at which learning is embedded within pricing decisions: predictive support, predict-then-optimize frameworks, and adaptive policy learning; and examine how these approaches shift modeling assumptions and structural integration in ancillary revenue management.

2. AI Methodologies in Ancillary Pricing

To structure the emerging literature on AI-based ancillary pricing, we classify contributions according to the degree to which learning is embedded within pricing decisions. We distinguish three stages of methodological integration: predictive support, predict–then–optimize frameworks, and adaptive policy learning.

In Stage 1 (predictive support), AI is used to estimate demand-related quantities such as purchase probabilities or willingness-to-pay. Learning improves demand prediction, but pricing remains externally determined. In Stage 2 (predict-then-optimize), AI-generated forecasts are embedded within an optimization framework. Demand is first estimated using machine learning, and prices are subsequently determined through a separate optimization layer. In Stage 3 (adaptive policy learning), learning and pricing are unified within a single decision process. Algorithms select prices while simultaneously updating beliefs about demand from observed outcomes. These stages represent increasing levels of decision autonomy and structural integration in ancillary revenue management systems. Figure 1 positions recent contributions within this framework, highlighting both methodological progression and areas of sparsity. We examine each stage in turn.

Article in Press

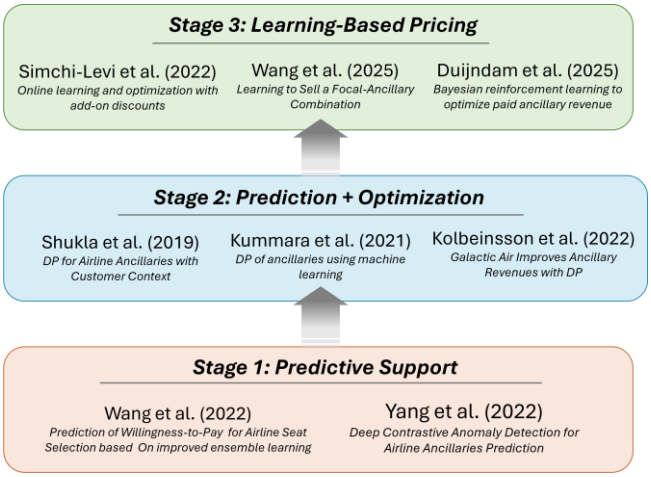


Figure 1. Stages of AI integration in ancillary pricing research. DP = Dynamic Pricing.

2.1 Predictive Support: AI for Demand Estimation

The first group of studies applies AI primarily as a predictive and computational support tool for ancillary pricing. In these contributions, machine learning improves demand estimation or computational tractability, but pricing decisions remain separated from the learning process. Wang et al. (2022) use LightGBM to predict willingness to pay for seat selection, generating purchase propensity estimates that inform pricing but do not determine prices directly. Yang et al. (2022) propose a deep contrastive learning framework to address severe class imbalance in ancillary purchase prediction, improving discrimination between purchasers and non-purchasers. In both cases, AI enhances demand estimation while pricing remains external.

2.2 Predict-Then-Optimize Frameworks for Ancillary Pricing

A second set of papers integrates AI predictions more directly into pricing decisions through a predict-then-optimize structure. Shukla et al. (2019) develop context-based ancillary pricing models that incorporate rich customer session information. They compare two-stage and end-to-end neural approaches that map estimated purchase probabilities to prices using parametric or search-based mechanisms. Although machine learning informs pricing, estimation and optimization remain distinct components.

Kolbeinsson et al. (2022) present large-scale deployment of personalized ancillary pricing using supervised learning and A/B testing. Pricing decisions are selected based on predicted demand, with periodic retraining for operational updating. However, pricing is not framed as an exploration–exploitation problem, and capacity and primary–ancillary interactions are not embedded in the learning framework.

Article in Press

Kummara et al. (2021) describe an implementation of machine-learning-based ancillary pricing using gradient boosting to estimate purchase probabilities. The approach integrates these predictions into an automated pricing decision process within a merchandising system. While the system enables more responsive and data-driven ancillary pricing, learning and pricing remain separated in a predict–then–optimize structure, ancillaries are treated independently of the primary product, and capacity interactions are not modeled explicitly.

Across these studies, AI enhances demand estimation, computational tractability, or price personalization, yet pricing logic remains separated from learning.

2.3. Adaptive Policy Learning for Ancillary Pricing

The third stage embeds learning directly within the pricing policy. Rather than estimating demand and solving a separate optimization problem, pricing is treated as a sequential decision process in which actions are updated based on observed outcomes. Performance is typically evaluated through regret, cumulative revenue loss relative to an ideal policy that knows the true demand model in advance (Bubeck & Cesa-Bianchi, 2012), or through simulation benchmarks. Despite this methodological advance, integrated learning in ancillary pricing remains limited and structurally simplified. The few contributions at this stage pursue distinct modeling directions, which we examine next.

Simchi-Levi et al. (2022) study revenue management with add-on discounts, where discounts on selected supportive products are offered conditional on a customer purchasing a core product. The retailer chooses prices for core and supportive products and selects which supportive products receive add-on discounts, subject to a space constraint limiting the number of displayed discount offers. In the offline setting with known demand functions, the problem becomes a mixed-integer/combinatorial optimization problem, and the authors develop a fully polynomial-time approximation scheme (FPTAS) that achieves any desired accuracy with polynomial runtime in the problem size. Because demand parameters are unknown in practice, this optimization routine is embedded within an online learning framework. They propose a UCB-based learning algorithm that constructs confidence bounds on demand parameters and uses the FPTAS as an optimization subroutine to compute an optimistic policy, updating decisions over time. The framework accommodates multiple add-ons under display constraints and provides regret guarantees for the online learning algorithm. While it advances adaptive pricing in combinatorial environments, it abstracts from capacity constraints and airline-specific primary–ancillary interactions central to ancillary revenue management.

Duijndam et al. (2025) develop a capacity-aware adaptive pricing model for a single paid ancillary (e.g., preferred seats) using Bayesian reinforcement learning. Demand is specified through an exponential function of price and approximated by a logistic-like sales function to incorporate the ancillary seat capacity constraint. The seller sets one fixed price per flight and

Article in Press

observes total seats sold at departure. Because demand parameters are unknown, they are treated as random variables with normal priors and updated in a Bayesian manner after each flight. Pricing decisions are generated via Thompson sampling: parameters are drawn from the posterior distribution, and the price that maximizes expected revenue under those sampled parameters is selected. Learning occurs across flights through a hierarchical prior that enables information sharing while allowing route-level heterogeneity. Although the primary ticket is treated as exogenous and only one ancillary is considered, this study provides one of the clearest integrations of Bayesian learning and explicit inventory control in ancillary pricing.

Wang et al. (2025) study learning and pricing for a primary-ancillary combination sold in sequence, an essential focal product is offered first, and only conditional on purchase is a discretionary ancillary offered. This purchase order creates an information dependency, ancillary valuations are only observed when the focal product is purchased, so focal pricing affects not just revenue but also the rate of learning about the ancillary (a censored problem for the ancillary). The paper formulates the learning problem in a bandit-style learning framework with customer covariates, designing policies with sublinear regret despite the censored feedback structure. They analyze three operational scenarios: (i) pure unbundling (separate focal and ancillary prices), (ii) a personalized mechanism that chooses whether to bundle or sell sequentially based on customer covariates, and (iii) a one-switch policy that allows only a single strategic switch over the horizon. For each case, the authors design learning-and-pricing algorithms with sublinear regret, establishing that the policy improves with experience despite the sequential dependence. Relative to Simchi-Levi et al. (2022), which learns discounts/assortment decisions over multiple add-ons under display constraints, and Duijndam et al. (2025), which learns prices under explicit capacity for a single ancillary, Wang et al. (2025) considers a single ancillary product and focuses on the learning challenges induced by the primary–ancillary purchase sequence.

Collectively, these studies demonstrate that adaptive learning has entered ancillary pricing through distinct modeling pathways, each isolating a different structural feature of the problem. The three-stage framework should remain a valid organizing lens as the literature grows, though Stage 2 may broaden as new hybrid approaches emerge. The stages are best understood as degrees of AI integration rather than discrete categories.

3. Assumption Shifts in AI-Based Ancillary Pricing

The progression from predictive support tools to adaptive policy learning reflects not only methodological evolution, but a redistribution of modeling assumptions. Classical revenue management approaches rely on explicitly specified demand functions and tractable optimization structures (Gallego & Topaloglu, 2019). AI-based approaches relax explicit

Article in Press

functional assumptions but introduce new assumptions related to data, experimentation, and environmental stability.

3.1 Demand and Data Assumptions

In predictive and hybrid frameworks, parametric demand curves are replaced with flexible function approximators such as gradient boosting and neural networks. While this reduces the need to pre-specify elasticities, it introduces implicit assumptions regarding data quality and stability. These models require sufficient historical price variation, stable feature distributions, and behavioral stationarity for reliable generalization (Hastie et al., 2009). In adaptive policy models, uncertainty is formalized differently but remains structurally constrained. Simchi-Levi et al. (2022) and Wang et al. (2025) assume unknown but fixed demand parameters and derive regret guarantees under repeated interaction. Duijndam et al. (2025) assumes a parametric demand structure and updates beliefs via Bayesian inference across flights. Despite the models' differences, all models rely on a stationary demand process during the learning horizon.

3.2 Experimentation and Learning Infrastructure

Adaptive pricing models introduce operational assumptions often understated in theoretical analysis. Contextual bandits, UCB-based algorithms, and Bayesian reinforcement learning require repeated interactions, observable feedback, and controlled experimentation (Cesa-Bianchi & Orabona, 2021; Bubeck & Cesa-Bianchi, 2012). Regret guarantees assume sufficient traffic volume and environmental stability for convergence. Even in deployment-oriented systems such as Kolbeinsson et al. (2022) depend on experimentation infrastructure. A/B testing and systematic price variation are necessary to overcome historical price rigidity and enable reliable estimation (Siroker & Koomen, 2015). Thus, AI-based ancillary pricing presumes organizational readiness for experimentation and continuous model updating.

3.3 Persistent Structural Simplifications

Despite advances in learning methodology, structural simplifications persist. The primary product is frequently treated as exogenous, and ancillaries are often priced independently. Capacity constraints are rarely integrated with adaptive learning, appearing explicitly only in Duijndam et al. (2025). Multi-ancillary combinatorial interactions are addressed in Simchi-Levi et al. (2022), but without capacity coupling or full primary-ancillary integration. Wang et al. (2025) explicitly models primary-ancillary dependence, but without inventory constraints. The literature therefore reflects methodological progress alongside structural compartmentalization. AI enables adaptive estimation and policy learning, yet no existing framework simultaneously integrates primary-ancillary dependence, capacity constraints, and multi-ancillary combinatorial complexity. Table 1 summarizes how assumptions shift across the three stages of AI integration in the reviewed literature.

Article in Press

Table 1. Key assumptions across AI-based ancillary pricing integration stages.

Assumption Dimension	Stage 1: Predictive Support	Stage 2: Predict-Then-Optimize	Stage 3: Adaptive Policy Learning
Demand Specification	Implicit, non-parametric function approximation	Forecasted demand embedded	Parametric or policy-level learning
Stationarity	Required	Required	Required
Price Experimentation	Not Inherent	Often via A/B testing	Endogenous to algorithm
Primary-Ancillary Integration	Typically absent	Rare and stylized	Explicit only in Wang et al. (2025)
Capacity Constraints	Absent	Rare	Explicit in only Duijndam et al. (2025)
Multiple Ancillaries	Independent modeling	Limited	Limited

4. Future Directions in AI-Based Ancillary Pricing

AI-enabled dynamic pricing has expanded rapidly across industries. Particularly, reinforcement learning and bandit frameworks are increasingly studied in pricing environments (Chenavaz & Dimitrov, 2025; Lange et al., 2025). However, their application of these methods within ancillary revenue management remains limited. Three gaps are evident. First, joint primary-ancillary pricing remains largely unaddressed. Second, adaptive learning models are rarely embedded within realistic capacity-constrained environments. Third, multi-ancillary combinatorial complexity has yet to be meaningfully integrated into AI-driven pricing systems. No existing study integrates these elements simultaneously. Beyond research gaps, practical and ethical considerations merit acknowledgment. Deployment requires clean transactional data, computational infrastructure, and organizational readiness for continuous retraining, while personalized pricing based on behavioral signals raises fairness and regulatory concerns. Future work should examine how such constraints can be embedded within AI-based ancillary pricing.

References

Bajari, P., Nekipelov, D., Ryan, S. P., & Yang, M. (2015). Machine learning methods for demand estimation. *American Economic Review*, *105*(5), 481-485.

Bubeck, S., & Cesa-Bianchi, N. (2012). Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *arXiv preprint arXiv:1204.5721*.

Cesa-Bianchi, N., & Orabona, F. (2021). Online learning algorithms. *Annual review of statistics and its application*, *8*(1), 165-190.

Chenavaz, R. Y., & Dimitrov, S. (2025). Artificial intelligence and dynamic pricing: a systematic literature review. *Journal of Applied Economics*, *28*(1), 2466140.

Article in Press

Duijndam, K., Koole, G., & van der Mei, R. (2025). Bayesian reinforcement learning to optimize paid ancillary revenue in the airline industry: K. Duijndam et al. *Journal of Revenue and Pricing Management*, *24*(6), 551-567.

Gallego, G., & Topaloglu, H. (2019). *Revenue management and pricing analytics* (Vol. 209). New York: Springer.

Hastie, T., Tibshirani, R., Friedman, J. H., & Friedman, J. H. (2009). *The elements of statistical learning: data mining, inference, and prediction* (Vol. 2, pp. 1-758). New York: springer.

Kummara, M. R., Guntreddy, B. R., Vega, I. G., & Tai, Y. H. (2021). Dynamic pricing of ancillaries using machine learning: one step closer to full offer optimization. *Journal of Revenue and Pricing Management*, *20*(6), 646-653.

Lange, F., Dreessen, L., & Schlosser, R. (2025). Reinforcement learning versus data-driven dynamic programming: a comparison for finite horizon dynamic pricing markets: F. Lange et al. *Journal of Revenue and Pricing Management*, 1-17.

Mumbower, S. (2025). Emerging research on airline ancillaries: a review of offer management and dynamic pricing literature. *Journal of Revenue and Pricing Management*, *24*(4), 353-367.

Ødegaard, F., & Wilson, J. G. (2016). Dynamic pricing of primary products and ancillary services. *European Journal of Operational Research*, *251*(2), 586-599.

Peng, Y., Choi, B., & Xu, J. (2021). Graph learning for combinatorial optimization: a survey of state-of-the-art. *Data Science and Engineering*, *6*(2), 119-141.

Siroker, D., & Koomen, P. (2015). *A/B testing: The most powerful way to turn clicks into customers*. John Wiley & Sons.

Shukla, N., Kolbeinsson, A., Otwell, K., Marla, L., & Yellepeddi, K. (2019, July). Dynamic pricing for airline ancillaries with customer context. In *Proceedings of the 25th ACM SIGKDD International Conference on knowledge discovery & data mining* (pp. 2174-2182).

Wang, H., Li, X., & Talluri, K. (2025). Learning to Sell a Focal-Ancillary Combination. *Naval Research Logistics (NRL)*, *72*(8), 1117-1138.

Wang, Z., Han, X., Chen, Y., Ye, X., Hu, K., & Yu, D. (2022). Prediction of willingness to pay for airline seat selection based on improved ensemble learning. *Aerospace*, *9*(2), 47.

Wilson, J. (2025). Analytical Modeling of Ancillary Items. *Analytics*, *4*(2), 14.

Wilson, J. G., & Ahmed, F. F. (2024). Pricing optimization for perishable primary and ancillary items with time-dependent inventory drawdown. *IFAC-PapersOnLine*, *58*(19), 592-597.

Vinod, B. (2023). Artificial Intelligence in travel. In *Artificial Intelligence and Machine Learning in the Travel Industry: Simplifying Complex Decision Making* (pp. 163-170). Cham: Springer Nature Switzerland.

Wilson, J. G. &Anderson, C. (2026). The Impact and Design of Slider-Based Offer Mechanisms in Mobile Commerce. To appear in *European Journal of Operational Research.*

Yang, P., Kolbeinsson, A., Shukla, N., & Barria, J. A. (2022, December). Deep Contrastive Anomaly Detection for Airline Ancillaries Prediction. In *2022 21st IEEE International Conference on Machine Learning and Applications (ICMLA)* (pp. 1167-1174). IEEE.

Yeoman, I. (2025). Expanding the frontiers of revenue and pricing management. *Journal of Revenue and Pricing Management*, 1-3.

Article in Press