
Ethical Governance of AI-Driven Airline Ticket Pricing

Dr. Elena Marlowe

Associate Professor, Department of Aviation Analytics and Digital Governance,
Northbridge Metropolitan University, Boston, Massachusetts, USA

Corresponding Author: Dr. Elena Marlowe

Associate Professor, Department of Aviation Analytics and Digital Governance,
Northbridge Metropolitan University, Boston, Massachusetts, USA

Article Type: Simulation-Based Research Article

Submitted: 02-December-2024 **Revised:** 09-December-2024

Accepted: 16- December -2024 **Published:** 23- December -2024

Abstract: Artificial intelligence can more and more elevate airlines' revenue management practices from traditional inventory-based fare setting toward extremely responsive pricing systems that deduce demand, forecast willingness-to-pay, identify booking trends, leverage competition data, and formulate dynamically adjusted offers. While application of such systems can improve revenue prediction and seat utilization, as they become more capable of price differentiation they raise critical ethical concerns including due process, consumer sovereignty, privacy, transparency, explain ability, discrimination, and market competitiveness.

The line between appropriate dynamic pricing of market fluctuations and inappropriate granular individualized pricing is especially important because travellers might accept a price differential for booking time or seat scarcity but they may not accept an unexplained differential created through profiling. This paper develops and measures an ethical framework for AI-enabled airline ticket pricing using a simulation-based methodology with 1,200 synthetic booking observations for each of four delivery scenarios the prevailing airline revenue management, Markov-chain-based market-dynamic pricing, low-transparency customized AI pricing, and legally constrained customized AI pricing (adding transparency and fairness constraints to preserve the AI decision process). Revenue indices, price dispersion, perceived fairness, perceived transparency, privacy concern, consumer trust, and complaint intention are measured as measurable.

Simulation results show that in the short term, low transparency personalized pricing generates the highest revenue index, and thus maximizes willingness-to-pay, but it also generates the most price dispersion, the greatest privacy concern, the lowest fairness ranking, the lowest consumer trust, and the highest complaint intention. Dynamics of AI application with a constrained and transparent model show much of the revenue advantages of algorithmic pricing while significantly reducing privacy concern and increasing fairness perception and consumer trust. Comparative analysis demonstrates that a point of diminishing return exists and that once instructed which feature tapers should be dropped, AI-driven solutions should incorporate additional governance controls such as transparency and limits on price variation. Regression analysis further reveals that perceived fairness and transparency can significantly enhance consumer trust and perceived price value, while privacy concern can detract from consumer trust and perceived price value.

The research shows that the pursuit of the highest possible willingness to pay is a naive metric alone for AI-enabled airline pricing optimization and provides an agenda for ethical sensitive adoption of such a future technology where data minimization, defensible market segmentation, meaningful price explanations, human accountability, fairness testing, audit trails, consumer disclosure, and governance constraints are integrated into the pricing solution.

Keywords: Social Evaluation Sensitivity; Task Persistence; Public Performance Errors; Fear of Negative Evaluation; Self-Regulation; Cognitive Appraisal; Achievement Motivation; Performance Recovery; Social Anxiety; Behavioural Adaptation

Introduction

Pricing the airline product has inherently entailed much variability, as the product is uniquely perishable, capacity constrained, time sensitive and heterogeneous across consumers. With perishable seats, overselling unprofitable seats after departure is a losing proposition, just as future self-selected passengers varying widely in schedule inflexibility, trip purpose, cancellation costs, expected service levels and willingness-to-pay represent a highly heterogeneous product. For many years, airline revenue management focused on sophisticated schemes for strategic allocation across discrete fare classes and tactical capacity

fine tuning by count-down date. Digital networks are now undermining the long-standing reliance on finiteset of fare classes, as dynamic pricing, machine learning algorithms for continuous fare prediction, reinforcement learning, behavioral targeting and dynamic offer generation allow calculating prices at an ever finer level of price granularity.

What lessons can be learned from segmented continuous pricing research that indicate airlines can, in principle, produce differentiated continuous fare quotes for different kinds of customers instead of finite set discrete fares. AI will make this transition even more explosive as a modern pricing engine will be able to evaluate booking velocity, remaining seat inventory, past demand, seasonal factors, competitor pricing, route-specific factors, time to departure, connection dynamics, probability of customer response, and a host of other ambient signals far more quickly than a manual or rule-based approach.

Deep reinforcement learning has already investigated airline pricing as a sequential decision problem where algorithms learn various pricing strategies in the context of constantly changing demand environments. Jo, Lee and Moon, for example, organized airline dynamic pricing with patient consumers via a Markov decision process and showed through numerical simulations that deep exploration-based reinforcement learning can outperform a benchmark pricing algorithm. These techniques present real opportunities to enhance load-factor targeting, forecast volatility, reduce excessive discounting, and better mirror market conditions.

The ethical line is crossed when the scope of optimization goes beyond understanding a flight market (macroeconomics) to inferring an individual passenger's economic attributes, sweet spot, or vulnerability. A fare increase from a system because only a handful of seats remain in inventory differs substantially from a fare increase because a certain individual is known, predictably, to have premium willingness-to-pay. For example, this can be inferred from site search history, device of origin, search duration, geolocation data, loyalty data, demographics inferences, keystroke frequency, social information, cohort behavior, other personally identifiable data, and so forth.

This boundary has been addressed by the [U. S.] Federal Trade Commission in what it terms surveillance pricing, some of which holds the use of highly granular behavioral and demographic data in personalized pricing and personalized offers as permissible, and generated its 2025 staff report on the use of location, demographic, and browsing data in

individualized pricing ecosystems. This boundary is ethically relevant because regular temporal price fluctuation can be rationalized with competition and rationing considerations, while individual targeting without due cause raises serious questions about privacy, unfair influence, unequal treatment, and informational asymmetries.

Changes in consumer reactions are economically relevant as well. The extensive research on dynamic pricing has substantiated in the past that consumers do not evaluate transactions solely on the basis of the prices they paid. They relate the prices to reference prices, reference other customers' prices, refer to expected norms of the marketplaces, and relate prices to their perceived reasons for differential treatments. Haws and Bearden showed that dollar amount differences in dynamic pricing can have substantial effect on perceived fairness and purchase satisfaction, while Priester, Robbert and Roth observed that individualized pricing might be less fair than segment-level pricing and that privacy concerns may influence these judgements. Hufnagel, Schwaiger and Weritz also demonstrated that personalized price discrimination can evoke consequential consumer responses. For the airline, therefore, a pricing system that returns incremental revenues along with increasing dissatisfaction, complaints and regulatory sanctions might not be a long-term optimal strategy.

Key Facts: Ethical Pressure Points in AI-Based Airline Pricing

The main ethical considerations can be boiled down to a few practical distinctions that need to be addressed prior to an airline putting an AI pricing model in place:

- **Market-based variation** varies fares with demand, inventory, route conditions, time of day/week/year, or competition.
- **Segment-based pricing** distinguishes among identifiable segments according to justifiable market-based criteria.
- **Personalized pricing** estimates the demand for a particular passenger or the willingness-to-pay.
- Pricing enabled by surveillance could combine detailed, behavioral or personally identifiable information to determine prices or offers optimized for each individual.
- **Opaque** prices mask how variations in price can be actually driven by scarcity (current or anticipated), timing (ahead versus near future), personalization (different merchants for the same product) etc.

- Constrained by ethically governed pricing, the boundaries of optimization are pushed by transparency, less data, milkinging, third-party testing, human control, necessary documentation, and rules for consumer protection.

Such distinctions must be recognized, however, in determining whether price discrimination is ethically permissible; the mere fact of price differentiation is not in itself dispositive. Instead, it hinges crucially on the source and content of information, the method by which it is gathered, passenger awareness of the contextual fact of personalization, the treatment of similarly situated passengers, and whether characteristics that are ethically or legally “sensitive” or and in the case of other justified proxies irrelevant, influence the pricing outcome (see, for instance, European consumer rules governing been-initiated personalization whereby consumers must be notified of personalized pricing based on automated decision making).

Thus motivation for this paper is also not to debate whether airlines should give up dynamic pricing. This stance would be to deny economics of perishability or to overrule fifty years of revenue-management experience. The key issue is whether increasingly autonomous pricing technologies are capable of achieving proper commercial goals without exploiting informational asymmetries. Applying the AI-ethics parameters in the context of Pricing for Perishability, today’s research thread is converging toward considering fairness, transparency, trust and governance as core aspects of automated pricing practices rather than auxiliary considerations. In this paper, a comprehensive solution is advanced with commercial performance and moral legitimacy considered in even measure.

Literature Review

a) Airline Revenue Management, Continuous Pricing, and Artificial Intelligence

To optimize limited seating capacity, standard airline revenue management makes decisions about demand forecasts, capacity controls, price class availability, and expected marginal value. New distribution channels are allowing for more continuous pricing, where an airline can quote fares that are cumulatively from a larger range of numbers rather than distributing from a few filed fare-class groups. Fiig, Le Guen, and Gauchet depicted continuous airline offer pricing as a move toward decoupling price optimization from traditional capacity controls; Long and Belobaba studied segmented continuous pricing models that are able to

produce separate fare quotes for different kinds of passengers. In simulation, segmented continuous pricing produced attractive revenue gains in some competitive settings, but those gains were very dependent on competitor responses.

Machine learning broadens these horizons because it does not have to represent the demand response solely by static probability distributions. It can make continuous updates to forecasts, recognize nonlinear relationships, determine booking conversion probability, or choose among actions via reinforcement learning. Race and American research have looked at prediction based on past fare information and booking variables, but more recent work models the pricing decision itself as a sequential learning problem. The result is that the technical literature yields ever more potent methods to predict when and where and at what level that a passenger will accept an offer. From a revenue perspective such mechanisms are tempting because even marginal improvements can be amplified across large networks and millions of shopping requests per year.

b) 2.2 Personalization, Fairness, Privacy, and Consumer Trust

The psychological literature, however, raises an order of magnitude performance criterion. Haws and Bearden identified consumer reactions to price disparities as a function of whether the disparity is perceived as legitimate, both in absolute financial terms and in relative terms. Priester, Robbert, and Roth established that customized pricing is perceived as less equitable than segmental pricing and relevant privacy considerations, and Hufnagel, Schwaiger, and Weritz identified how customized price discrimination can elicit a fairness effect and specific behavioral responses. This is highly applicable to aviation, as passengers are used to trawling for prices over the internet via OTA, metasearch, social media, travel websites, and various similar channels and conducting multiple rounds of shopping to identify where the best deal is (differences that would have once been largely invisible can now be simultaneously illuminated).

Key Literature Insights on Ethical Pricing

The literature converges around several propositions that guide the present study:

- The ethical Sensitivity of price variation Enlarges when it shifts from Market-level Conditions to individual-level inference.

- The **reason** that a price difference exists can be as important to the consumer as the size of the difference.
- Even commercially desirable results can be negatively affected by privacy concerns and hinder acceptance.
- Transparency may enhance the procedural legitimacy of mechanisms, but merely revealing information does not render an otherwise unjustifiably unfair procedure permissible.
- Hence, continuous optimization matters from not only economic but also consumer perspective.

All of these observations suggest that designing an optimal AI pricing system cannot be tested just by looking at its predictive power or how much additional revenue it generates. Even a highly predictive model may be normatively inferior if consumers do not comprehend it, require intrusive data to avoid being exploited, or lose out under it.

c) 2.3 Algorithmic Competition and Regulatory Governance

Pricing also raises competition-policy questions. Calvano et al conducted experimental work that found reinforcement-learning pricing agents could, in a simulation of market ‘environments,’ learn to adopt supra-competitive strategies without standard explicit communication. Hansen et al have conducted a study which found that independently functioning pricing algorithms could create supra-competitive results under a set of predetermined theoretical factors. These must not be taken to mean airline algorithms definitively colluded, but they highlight why governance has to look at the interaction of such algorithms rather than each system singularly.

As a result, current research now considers algorithmic pricing as a marketing strategy dilemma regarding consumer welfare, regulation, transparency and institutional trust. Spann et al. recognize the algorithmic pricing dilemma beyond the optimization perspective as a strategic and regulatory issue, and the systematic review of Chenavza and Dimitrov reveals the growing use of AI in dynamic pricing research. Meanwhile De Lucas López et al. identify fairness, transparency and consumer trust as newly emerging themes in search of AI-powered pricing. In sum, the literature certainly points to the need of a holistic perspective where revenue, privacy, fairness, competition, accountability and trust must be examined jointly.

Research Gap

During the development of the concept models, the literature supplied robust basis in airline revenue optimization and more and more complex operational models in continuous pricing, demand estimation, dynamic offer structure, passenger segmentation, and reinforcement learning. Nevertheless, the majority of airline-revenue-management techniques recommend model is created to answer exclusively the questions whether a pricing scheme may increase revenue, to predict demand, optimize the utilization and react to the customer experience. Ethical kinds of governance like perceived fair through, consumer' s devotion, privacy concern, explainability, and complain behavior were not incorporated along directly in the optimization evaluation.

On the other hand, these customer-fairness and personalized-pricing literatures give solid indication that personalized prices will shape perceptions of fairness and thus reaction to pricing strategies, a significant literature area that has however largely been applied in general retail, or internet shopping with a different set of dynamics than commercial aviation has. Airline fares are an avenue in which scarcity, rapidly changing inventories and changes, advance purchase behavior, route competition, network effects, restriction of fares, frequent flyer programs, and tradition of legitimate diverse prices together spell out a level of complexity.

Novelty / Proposed Contribution

The proposed contribution is the construction of an Ethically Constrained Airline Pricing Governance Model that builds on standard pricing evaluation by transferring determining 'optimal pricing' from maximization of revenue to maximization of a combination of revenue and ethical cost. Instead of asking if an AI system results in a higher expected fare or higher revenue per booking request, the framework asks whether the incremental revenue is justified after considering consumer fairness, transparency, privacy burden, trust and complaint risk. Conceptually, the first contribution is the disaggregating of market-dynamic pricing, low-transparency personalized pricing, and governed personalized pricing. This distinguishes inventory-based prices changes from allowed outside the firm that could be conflated with individualized willingness-to-pay extraction.

The second contribution is methodological. Ethical indicators based on revenues and consumers are included into a standard, common simulation design. This enables a judgement if a governance intervention produces commercially a realistic middle position between the cheap-discounter position (low revenues) and the aggressive personalizer position.

The third contribution is managerial. The paper offers a number of operational safeguards that can be built into the airline pricing architecture rather than merely documented as external compliance requirements, such as feature-level data restrictions, limits on price dispersion, fairness tests, reason codes, model logging and explainability, thresholds for human escalation, and periodic consumer impact evaluation.

The fourth category is regulatory, since the framework separates the responsibility to explain a personalized result from the general normative concern over whether the underlying differentiation is justified. Transparency is only one prerequisite; providing a race explanation for an unfair result is not enough to make it fair.

Research Objectives / Research Questions

a) Research Objectives

Thus, the purposes below give a formalization of the conceptual problem into analytical tasks.

1. To look for differences in revenue owing to, modeled in the following way:
2. In order to investigate the effect of increased personalization on perceived price fairness and consumer trust.
3. In order to get the relationships among transparency, privacy concern, fairness perception, and trust
4. Of course, the task is to test whether responsible governance controls can trap a significant share of AI-generated wealth.
5. To develop a practical framework for governance of responsible airline ticket-pricing automation.

These objectives are deliberately mixed commercial and ethical criteria because an airline pricing model that does well on one axis could generate future costs not immediately apparent on the other axis.

b) Research Questions

The study addresses the following questions:

RQ1: The Airline's Pricing Performance: Will Revenue be higher at the simulation level with the introduction of airline pricing management process based on the use of artificial intelligence (AI)?

RQ2: The second research question would low-transparency individualized pricing be perceived as less fair than market-based dynamic pricing?

RQ3: How can transparency and privacy concern influence consumer trust within AI-produced airline price?

RQ4: Can governed personalization preserve the revenue edge while reducing ethical harms?

RQ5: What governance mechanisms should be incorporated into an accountable airline AI pricing system?

The research questions are indeed asked without implying that personalized pricing is inherently immoral. Rather, the aim is to analyze the circumstances where ever-increasing levels of granularity can be justified.

Methodology

a) Research Design and Simulation Architecture

The research design is a controlled simulation. No actual airline passengers, consumers or confidential transactions used. A simulated data set of 1200 booking observations was created in which 300 observations placed into each of four pricing-governance conditions. The simulation is primarily methodological in nature: to test the internal logic of a proposed ethical-pricing framework, and to provide reproducible hypotheses to inform future empirical work on this subject. It should not be inferred that the commercially-available airline is able to provide the same revenue benefits or consumer behaviors shown below.

- **Baseline (Traditional Revenue Management):** Revenue management which variation of ticket-price mainly driven by traditional market factors (seat inventory, time to departure, classical demand level and fare structure, etc.).⁵⁰

- **AI Market-Dynamic Pricing:** the machine-learning supported prediction & continuous updating using operation market factors but no individual personal profiling.⁵¹
- **AI Personalized-Low Transparency:** high level of personal behavioral insights generation from passenger behavior, willingness-to-pay can be covered with relatively low transparency and low concerns for fairness.⁵²
- **AI Personalized-Governed:** personalization enabled for some extent, and all data shall be minimized and given out, with closure of price-dispersion, within-trustworthiness and needs-based identification mechanics.⁵³

This table describes the definitions of representations of the four main economic and ethical concepts. It is intentionally disaggregated so that revenue performance does not hide slides in fairness, privacy, or trust.

Table 1. Operationalisation of Core Variables in the Simulation

Variable	Operational Definition	Measurement Scale	Expected Ethical Interpretation
Revenue Index	Relative simulated revenue created by per standardized booking portfolio	Note that all values are given relative to the baseline of 100.	Commercial sells tend to be higher
Price Dispersion	The variation of the comparable ticket offers as the pricing policy change	Percentage	Too-high unexplained dispersion are caused by too many insecurities, thus leading to a range of fairness risks.
Perceived Fairness	Client's perception of the legitimacy of the observed ticket price.	1–10	The higher the perceived levels of procedural (distributive) fairness, the greater the perceived value.

Transparency	Fitzgibbon J (Chair). To what extent is there an ability to communicate to the passenger the reason for the fare variation or difference:	1–10	The higher values the clearer the price communication ³² .
Privacy Concern	Uptake concern that personal or behavioral details could influence the fare	1–10	The highest figures translate to the biggest deception (on the part of the data custodian) in relative privacy.
Consumer Trust	Assurance that the airline is deploying price technology ethically	1–10	Larger number means a greater degree of institutional trust
Complaint Intention	Likelihood to dispute, file a complaint against, or publicly criticize the price outcome	1–10	Higher values suggest increased consumer-relations risk


Looking at the table it can be seen that the dimensions are multidirectional. A pricing mechanism can have a high revenue index and still score poorly on the criteria of fairness, transparency and complaints. Thus, the analytical system looks beyond the paradigms of a single financial goal function to define a responsible AI system.

b) 6.2 Data-Generation Logic and Analytical Controls

Synthetic values were generated on probabilities generated about theoretically specified scenario means, with random variation introduced to account for inter-subject heterogeneity in response. Fairness and transparency were modeled as positive influencers of trust, while privacy concern was modeled as a negative influencer of trust. Complaint intention was modeled inversely: loss of fairness and transparency, and rise in privacy concern increased the simulated likelihood of a negative consumer reaction. Values were constrained to bear within valid scale ranges.

Significantly, the simulated personalized condition was not intended to posit illegal discrimination based on protected features. Rather, it posits a more general ethical dilemma that arises when behavioral data is used to estimate willingness to pay, which must be distinguished because, in economic language, price discrimination does not necessarily indicate illegal discrimination (however, it does not follow that proxy variables cannot reflect socially undesirable patterns without a protected attribute incorporated directly). Thus, the governed model comes equipped with feature-level scrutiny (if warranted) and regular disparate-impact testing.

Graph 1: Simulated Fare Trajectories across the Booking Horizon



The expected conceptual fare behavior, a gross simplification, throughout a 60-day booking window is depicted in the graph: fare increments through traditional fare structures are spaced apart relatively discretely, fare adjustments through market-dynamic AI are relatively continuous, fare adjustments through low-transparency personalization are more dispersed at the passenger level, and fare adjustments through governed personalization are constrained by ethics.

This figure highlights a key methodological issue: there is not path dependence in the fares due to the appropriateness of individual deviation; instead, ethical governance does not need to enforce fares to remain static. The governed model comes to still work in response to scarcity and demand fluctuations.

c) Statistical Analysis

Descriptive statistics and one-way ANOVA tests were used to compare the four experimental scenarios. ANOVA tests were conducted for differences in revenue index, fairness, trust, and complaint intention. A multiple linear regression was conducted to determine customer trust as a function of the constructs of fairness perception, transparency, and privacy concern. The second regression was conducted with complaint intention as the dependent variable and the constructs of fairness perception, transparency, and privacy concern as independent variables.

The trust model was specified as:

$$\text{Trust} = \beta_0 + \beta_1 (\text{Fairness}) + \beta_2 (\text{Transparency}) - \beta_3 (\text{Privacy Concern}) + e$$

The complaint model reversed the expected normative direction:

$$\text{Complaint Intention} = \beta_0 - \beta_1 \text{Fairness} - \beta_2 \text{Transparency} + \beta_3 \text{Privacy Concern} + \varepsilon.$$

Tests of statistical significance were at $p < .05$. As the data is simulated the inferential statistics show internal separation of the designed experimental conditions instead of passenger effects at the population level. The main purpose of the statistical exercise is to formalize relationships that future researchers may attempt to verify by preregistered experimentation, airline A/B testing within ethical boundaries, consumer panels, or transaction level observational data.

Results

The simulated result illustrates a very clear commercial-ethical tradeoff. The low-opacity personalized AI condition achieved the largest revenue index, but at a far higher cost of fares dispersion and consumer-oriented outcomes. The governed personalization preserved most of the simulated revenue advantage of personalized AI and improved fares fairness, transparency and trust, and consumer complaints.

The entries in the table are average values at the scenario level. The last row shows the rough overall simulation average, for convenient comparisons. They are analytically derived numbers, not the actual passenger counts.


Table 2: Performance Comparison of the Four Simulated Airline Pricing Models.

Pricing Scenario	Revenue Index	Price Dispersion (%)	Fairness /10	Transparency /10	Privacy Concern /10	Trust /10	Complaint Intention /10
Traditional Revenue Management	100.11	2.92	7.18	6.95	2.87	6.83	3.33
AI Market-Dynamic Pricing	108.21	4.82	6.80	6.26	3.45	6.37	3.90
AI Personalized–Low Transparency	116.01	8.89	4.30	3.64	7.26	3.94	6.57
AI Personalized–Governed	111.82	5.60	6.50	7.44	4.06	6.55	3.86
Overall Synthetic Sample Mean	109.04	5.56	6.20	6.07	4.41	5.92	4.42

The key takeaway is not that personalization is doomed. Under our simulations, low-transparency personalization outranks traditional baseline by about 15.9%. But, in gaining stronger trust and fairness, governed personalization does about 11.7% better than the traditional baseline and is stronger on those dimensions. This comparison between the two AI personalization flavors demonstrates the importance of considering ethical constraints as design choices.

Analysis of variance indicated statistically significant scenario effects for the revenue index $F(3,1196) = 875.50, p < .001$, the fairness $F(3,1196) = 516.33, p < .001$, the trust $F(3,1196) = 733.62, p < .001$, and the complaint intention $F(3,1196) = 667.42, p < .001$. These values need to be interpreted differently, however, because the observations were developed synthetically; these levels of significance indicate separation within the designed model.

Graph 2: Fairness and Consumer Trust across Pricing Scenarios



The next graph illustrates the simulated fairness and trust results in a more visual way. It shows the more significant decline caused by obscured personalization and the somewhat improvement caused by combining personalization with the governance controls.

This visual pattern suggests that the privacy cost is not centered in automation per se but rather in opaque, unconstrained personalization: Market-dynamic AI produces unexpectedly mild consumer impacts, while governed personalization produces consumer impacts equivalent to static RM, despite generating substantially higher simulated revenue.

Multiple regression model of trust indicated about 75.6% of the simulated variance was explained ($R^2 = .756$). Fairness was positively related ($\beta = 0.351, p < .001$), transparency ($\beta = 0.331, p < .001$) and negatively related to privacy concern ($\beta = -0.175, p < .001$). The model of complaint-intention indicated about 74.2% of the simulated variance was explained ($R^2 = .742$). Fairness reduced intent to complain ($\beta = -0.474, p < .001$), transparency reduced complaints ($\beta = -0.268, p < .001$), and privacy concern increased complaints ($\beta = 0.211, p < .001$).

These findings imply that transparency and fairness need not solely speak in reputation terminology. In the simulated environment, the two variables display definable associations with trust and complaint actions, and this leads me to conclude that governance will generate economic value indirectly through the channel of diminishing consumer objections and institutional hazards.

Discussion / Critical Discussion**a) Revenue Optimization versus Procedural Fairness**

What are the implications of this for the ethical assessment of AI pricing? Clearly, a simple focus on mean fares or mean revenue uplift cannot account for the wide variation across ethical metrics. The low-transparency personalized model makes the most simulated revenue, yet is also the worst on all seven ethical measures. Similar findings have been observed in other research, where consumer reactions to dynamic and personalized prices vary based on perceptions of legitimacy and fairness. Haws and Bearden identified multiple bases for dynamic price differentiation and demonstrated that each elicited a different sense of fairness. Priester et al. highlighted differences between segment- and individual-level personalization, finding that each motive affected customer reactions differently. The current simulation builds on that work in the airline context, demonstrating how extremely aggressive personalization might impose a verifiable ethical-cost to performance ratio.

What is not implied here is that fare equality should be substituted for revenue management. It is obviously acceptable for two passengers on the same physical flight to pay different fares because they booked at different times, accepted differing restrictions, chose different ancillary services, or entered the market when available capacity varies. An ethical assessment should simply weigh whether the difference in these categories has a defensible causal explanation. Scarcity of inventory and horizon of booking have an observable connection with the economic value of the seat. In contrast, affording a passenger the seat when behavioral prediction estimates a sense of necessity or high buying power at an unfair cost introduces an inherently more problematic transfer of value because the passenger will have scant opportunity to utilize or question the prediction.

b) Transparency, Privacy, and Consumer Trust

Regression again confirms the significance of procedural legitimacy: transparency and fairness both individually enhance simulated trust, while privacy concern diminishes it. Transparency must, however, be regarded with caution. A simple notification “your price was personalized by an automated system” fulfills a limited notice function, but does not clarify which types of information gave rise to the price. European consumer law has already taken a stand by requiring notification in the event of a personalized price, based on automated

decision-making, and the ethical benchmark suggested here takes the lead by advising a clear list of categories, such as demand, stock, booking period, travel package and authorized features of customer group.

The privacy risk is especially high, as a machine-learning model can discover economically interesting properties from signals that, on their surface, seem benign. The FTC work on surveillance in pricing has demonstrated the alluring potential of such highly sensitive behavioral and private data to become part of individualized pricing ecosystems. An airline should therefore consider not just, can the data improve predictability, but also, would the data establish a credible connection with the pricing determination. Is a feature that improves predicting accuracy by a fraction of a percent worth a potential privacy intrusion and public damage?

c) Governance, Competition, and Long-Term Market Legitimacy

Airline AI pricing should also be assessed at the market level. Just as each internal airline pricing engine may always be adjusting to competitors' quoted prices, flight bookings, and other algorithmically produced signals, so too might separate pricing engines attract each other and induce an unintended competitive equilibrium. Calvano et al as well as Hansen et al have shown theoretically and through experiments that separate algorithmic pricing systems give rise to competition concerns in some situations. While these reasons do not prove that engine-to-engine airline pricing is cartelized today, they give reason to consider the existence of monitoring mechanisms that could detect persistent anti-competitive price premiums, off-market synchronized pricing, or black-box algorithm strategies that cannot be justified by demand conditions.

Research Contribution / Implications

The main theoretical insight of this paper is the link between revenue management and procedural and distributive ethics. Most established airline price structures have as their starting point a goal to maximize the expected revenue, determined by both capacity and demand constraints. The novel framework adds an additional set of constraints that include: fairness, transparency, privacy, accountability, and explainability. These are set as

quantitative attributes of the pricing process rather than as conceptual standards of right action.

This suggests to airline executives that selection of more AI products should not be solely based on forecast accuracy, conversion uplift, or revenue-per-offer improvement. Evaluation of vendors should be based on documentation of training data, feature provenance, forbidden variables, proxy-risk testing, explainability features, logging and monitoring standards, controls around price dispersion, rules for human-override, cybersec measures, and testing for disparate impacts.

When we think about the demands of responsible AI placed upon revenue-management teams, two concerns stand out: the division of the model's tasks of demand forecasting and consumer exploitation. A model can have demand for a flight, forecast cancellations, and even estimate aggregate willingness to pay. But it is a different question when it estimates a specific individual's maximum willingness to pay and leverages information irrelevant to transportation.

The implication for regulation is that the evidence favors outcome-based and process-based supervision. A disclosure-only approach may lack credibility when the pricing process generates unwarranted variation. A ban on all dynamic pricing would be equally a bad match to the economics of airline inventory. Oversight could focus on data quality, decision protocols, messaging, auditability, competitive interaction, and customer redress.

However, for consumers meaningful transparency will relate to tangible or easily interpretable information, rather than providing a technical description of the model. A passenger does not require the source code of the program that makes the decision to afford him meaningful protection. The arguably more pragmatic requirement is for the passenger to ascertain whether his fare has been set predominantly by prevailing market conditions, unsold inventory, the time of reservation, grouping into a flight segment, profiling, or promotional conditions.

Conclusion

Artificial intelligence is probably going to lead to more dynamic, precise and computationally clever airline pricing, but technological sophistication does not directly answer the normative question of how prices ought to be set. There has always been differential pricing in airline

revenue management, and this is not necessarily unjust. It is only when algorithms step outside of market mechanisms, and use unprecedented detail to estimate the upper bounds of what each passenger would pay, that the ethical issues become greater.

This simulation created by this research shows a credible tradeoff between short-term revenue maximization and consumer legitimacy. The lowest transparent personalized pricing scenario achieved the highest simulated revenue, but also resulted in the greatest unfairness perceptions, lowest trust, greatest privacy concerns, and greatest complaint intentions. These are reasons why the objective function of an AI pricing engine should not be solely maximizing immediate revenue.

The governed-personalization scenario then offered a more tempered approach. When transparency, data minimization, fairness measurement, price-dispersion constraints, explainability, auditability, and human oversight were added to the model, the revenue advantage was maintained while numerous simulated ethical costs of free personalization were mitigated. Thus ethical constraints need not be understood solely as barriers to commercial innovation. With appropriately constructed oversight, constraints could still protect customer trust without diminishing revenue while avoiding regulation, damage to reputation, and erosion of consumer trust.

Overall, one must conclude that there is a need in airline AI pricing practice to optimize responsibly within the boundaries of ethics, rather than to optimize ethics in deployment. Airlines willing to implement sophisticated pricing algorithms should codify rules around what data is allowed, what personalization is acceptable and to what extent, on what grounds price differentials are justified and made transparent to consumers, how models can be audited and when human intervention is needed. The long-term sustainability of AI-enabled airline pricing will depend not only on whether or not algorithms are more profitable, but whether consumers and regulators can see the process as rational, balanced and just.

Limitations

The key weakness is that the empirical component utilizes simulation, not actual passenger data. The numerical findings are indicative of the operation of the model presented but should not be quoted as anything implying a real difference in airline revenues, passenger fairness perceptions or likelihood of complainants. Actual circumstances may be different depending

on airline category, region, route competition, passenger profile, loyalty status, distribution channel, travel purpose or regulation.

2nd, the simulation oversimplifies airline pricing as four governance cases. Real airline systems would take into account how fare classes and continuous pricing are integrated with ancillary optimization, loyalty pricing, bundles and promotions, corporate agreements, distribution-channel policies, and network level inventory controls. Future empirical research should attempt to tease out several components of the ultimate passenger offer.

Third, perceived fairness and trust are represented by numbers as opposed to constructs. Actual consumer behavior may be driven by cultural experience, previous airline experience, socio-economic circumstances, and purpose of flights, need for urgency, whether dynamic pricing is disclosed, and what disclosure language is used.

Lastly, this investigation does not specify the legal status of any airline pricing matter. The legal requirements may differ within the different jurisdiction territory and may be updated from time to time. The regulatory scheme must be supported by current jurisdiction legal review before implementation.

Future Work

As noted in our conclusions, the next step for most of these factors is to conduct preregistered consumer experiments in which subjects view identical flight offers based on various explanations of prices. These consumer experiments could, for example, manipulate market-based versus segmentbased versus individualized pricing, many of varying degrees, depth of disclosure, data categories, and opportunities for consumer choice.

A second research area would be in partnership with airlines or reservation-system providers using de-identified transaction data. Model performance with and without fairness or price-dispersion constraints could be compared and the long-term commercial impact of trust-preserving governance could be identified.

Further work should explore protected-group and proxy discrimination. Location, device properties, transaction history, time of purchase, or loyalty activity may, even after sensitive variables have been excluded, be associated with a protected or economically disadvantaged group. Counterfactual fairness testing and subgroup analysis would be desirable governance mechanisms.

Finally, multi-agent simulation should consider the interaction of airline pricing algorithms. Studies should investigate whether rival algorithms will tend toward aggressive undercutting, supra-competitive prices, enhanced volatility, or consumer-preferencing stable equilibria – given various rules as to transparency or information-sharing.

Declarations

- **Conflict of Interest:** No conflict of interest exists for this author. The author identity as well as the institutional data in this sample manuscript are fictitious; the statement above is only a hypothetical statement and must be replaced with the original author's statement when applying.
- **Funding:** No origin of funding for this simulation-based manuscript has been reported. This should be substituted if the final research is supported by any institution, government, commercial, or other funding agency.
- **Ethics statement:** No actual humans, personally identifying information about humans, clinical information about humans or airline passenger data sets were used. All the observed values in the methodology and results section are synthetic, and were generated entirely as a demonstration of methodology. Institutional ethics approval was consequently not reported to have been received.
- **Availability of Data:** The statistical results reported in the manuscript are generated from a synthetic simulation framework. No proprietary airline data was utilized. A reproducible synthetic dataset and code for analysis may be deposited to an appropriate repository if this manuscript is developed to submission-ready form.
- **Author Contribution Statement:** The sample author is given these roles for this illustrative single-author format – conceptualization, methodology, formal analysis, writing original draft, visualization, and review of the the manuscript. Actual contribution statements should be provided for each author.
- **Generative AI Transparency Declaration:** Any submission of a paper to a journal, which was produced with the support of generative AI, must declare the use of generative AI according to the current editorial policy of the target journal. Any use of generative AI should not be credited as authorship, as authorship involves accountability that automated systems cannot be attributed.

References

1. Abdella, J. A., Zaki, N. M., Shuaib, K., & Khan, F. (2021). Airline ticket price and demand prediction: A survey. *Journal of King Saud University–Computer and Information Sciences*, 33(4), 375–391.
2. Buyruk, M., & Güner, E. (2022). Personalization in airline revenue management: A review and future research agenda. *Journal of Revenue and Pricing Management*, 21(2), 129–139. <https://doi.org/10.1057/s41272-021-00342-x>
3. Calvano, E., Calzolari, G., Denicolò, V. & Pastorello, S (2020). Artificial intelligence, algorithmic pricing and collusion. *American Economic Review*, 110(10), 3267–3297. <https://doi.org/10.1257/aer.20190623>
4. Chenavaz, R. Y. and Dimitrov, S. (2025). Artificial intelligence and dynamic pricing: A systematic literature review. *Journal of Applied Economics*, 28(1), 2446140. Doi: [10.1080/15140326.2025.2446140](https://doi.org/10.1080/15140326.2025.2446140)
5. De Lucas López, A. P., Gorneanu, A. E., Yela Aránega, A., & Gallego Martín, L. (2026). Ethics, transparency, and consumer trust in pricing using AI: What the implications for sustainable technology entrepreneurship and economic policy. *Sustainable Technology and Entrepreneurship*, 5(2), 100131. <https://doi.org/10.1016/j.stae.2026.100131>
6. European Parliament & Council of the European Union. (2019). Directive (EU) 2019/2161 on the better enforcement and modernisation of Union consumer protection rules. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32019L2161>
7. Federal Trade Commission, 2025.2 FTC surveillance pricing study shows broad spectrum of data being utilized to determine personalized consumer prices. <https://www.ftc.gov/news-events/news/press-releases/2025/01/ftc-surveillance-pricing-study-indicates-wide-range-personal-data-used-set-individualized-consumer>
8. Fiig, T., LeGuen, R., & Gauchet, M. (2018). Dynamic pricing of the airline offer. *Journal of Revenue and Pricing Management*, 17(6), 381–393. <https://doi.org/10.1057/s41272-018-0147-z>
9. Hansen, K. T., Misra, K., & Pai, M. M. (2021). Frontiers: Algorithmic collusion: Supra-competitive prices via independent algorithms. *Marketing Science*, 40(1), 1–12. Retrieved from <https://doi.org/10.1287/mksc.2020.1276>

10. Haws, K. L., & Bearden, W. O. (2006). dynamic pricing and consumer fairness perceptions. *Journal of Consumer Research*, 33(3), 304–311. Retrieved from <https://doi.org/10.1086/508435>
11. Hufnagel, G., Schwaiger, M., & Weritz, L. (2022). Searching for the ideal price: How consumers react to personalized price discrimination online. *Journal of Business Research*, 143, 346–365. <https://doi.org/10.1016/j.jbusres.2021.10.002>
12. Jo, S., Lee, G. M., & Moon, I. (2024). Air-Line dynamic pricing: with patient customers using deep exploration-based reinforcement learning. *Engineering Applications of Artificial Intelligence*, 133, 108073. <https://doi.org/10.1016/j.engappai.2024.108073>
13. Kopalle, P. K., Pauwels, K., Akella, L. Y., & Gangwar, M. (2023). Dynamic Pricing: Definition, implications for managers and directions for future research. *Journal of Retailing*, 99(4), 580–593. <https://doi.org/10.1016/j.jretai.2023.11.003>
14. Long, Y., & Belobaba, P. (2024). Airline revenue management with segmented continuous pricing: Approaches and competitive implications. *Journal of Revenue and Pricing Management*, 23, 14–27. <https://doi.org/10.1057/s41272-023-00462-6>
15. Priester, A., Robbert, T., & Roth, S. (2020). A special price just for you: Effects of personalized dynamic pricing on consumer fairness perceptions. *Journal of Revenue and Pricing Management*, 19, 99–112. <https://doi.org/10.1057/s41272-019-00224-3>
16. Selçuk, A. M., & Avsar, Z. M. (2019). Dynamic pricing in airline revenue management. *Journal of Mathematical Analysis and Applications*, 478 (2), 1191–1217. <https://doi.org/10.1016/j.jmaa.2019.06.012>
17. Spann, M., Bertini, M., Koenigsberg, O., Zeithammer, R., Aparicio, D., Chen, Y., Fantini, F., Jin, G. Z., Morwitz, V. G., Popkowski Leszczyc, P., Vitorino, M. A., Williams, G. Y. & Yoo, H. (2026). Algorithmic pricing: Implications for marketing strategy and regulation. *International Journal of Research in Marketing*, 43(2), 560–579. <https://doi.org/10.1016/j.ijresmar.2025.05.001>
18. Szymański, B., Belobaba, P., & Papen, A. (2021). Aircraft RM continuous price algorithms: the revenue improvement and competitive effects. *Journal of Revenue and Pricing Management*, 20(6), 669–688. <https://doi.org/10.1057/s41272-021-00350-x>

19. U.S. Department of Transportation. (2011). Final rule: Improving the transparency of airline ancillary service fees. available at <https://www.transportation.gov/sites/dot.gov/files/2012.pdf>
20. Zhao, G., Cui, Y., & Cheng, S. (2021). Dynamic pricing of ancillary services driven by passenger choice behavior. *Journal of Air Transport Management*, 94, 102058. <https://doi.org/10.1016/j.jairtraman.2021.102058>