AI Adoption, Market Outcomes, and Coordination Risks*

Qi Ge† Minhae Kim‡ Myongjin Kim§

March 2026

Abstract: This paper examines the impacts of artificial intelligence (AI) adoption in the U.S. airline industry on pricing, service quality, environmental externalities, and consumer welfare. We find that enterprise AI adoption improves on-time performance and local air quality while raising average fares and compressing price dispersion. For generative AI–enhanced pricing tools, the price effects are larger and extend across the fare distribution, but their direction and magnitude depend on the extent and pattern of diffusion across carriers. In particular, when rivals rely on the same algorithm vendor, fares tend to be higher throughout the distribution, especially at the lower end, whereas the diffusion of GenAI-enhanced pricing platforms is associated with lower fares across the distribution, with the effect again being more pronounced at the lower end. Overall, our findings suggest a key trade-off in the adoption of AI technologies between efficiency gains and coordination risks.

Keywords: AI, generative AI, algorithmic collusion, pricing, quality, AQI, coordination risks, airline industry
JEL codes: L13, L22, L41, L93, O33, Q53

*We are grateful to Russell Pittman, Benny Mantin, Christian Bontemps, Angela Bergantino, Tyler Ransom, Mu-Jung Yang, and Nicholas G. Rupp, as well as seminar participants at the University of Luxembourg, the University of Bari, and the University of Oklahoma, and conference participants at the SoAR conference, for their valuable feedback. We also thank Ali Hortaçsu, Steven Tadelis, and Federico Ciliberto. Authors are listed alphabetically and contribute equally. All errors are our own.

†Department of Economics, Vassar College, Poughkeepsie, NY 12604. Email: *qige@vassar.edu*.

‡Department of Economics, Oklahoma State University, Stillwater, OK 74074. Email: *minhae.kim@okstate.edu*.

§Corresponding author. Department of Economics, University of Oklahoma, Norman, OK 73019. Email: *mjkim@ou.edu*.

1 Introduction

Artificial intelligence (AI) is transforming firm operations across different industries, ranging from manufacturing and logistics to healthcare and transportation. The airline industry has been among the leading adopters of enterprise AI systems that apply data-driven algorithms to improve operations, automate processes, and support business decision-making. Major U.S. carriers, including American Airlines (AA), United Airlines (UA), and Delta Air Lines (DL), have steadily expanded their use of AI technologies to enhance predictive maintenance, flight operations and scheduling, baggage handling, and customer service.[1]

Beyond operations, AI is playing an increasing role in competition and pricing. In the airline industry, existing enterprise AI systems can analyze large volumes of data to dynamically adjust fares in response to demand fluctuations, booking patterns, and competitor prices. More recently, generative AI (GenAI) applications are being piloted to extend these capabilities by simulating market conditions and consumer responses to generate individualized fares in real time. For example, Delta Air Lines rolled out a generative AI pricing engine, referred to by the carrier as a "super analyst," that estimates each passenger's maximum willingness to pay and sets personalized ticket prices. The airline expected to apply this system to price about 20% of its flights by the end of 2025.[2]

The use of AI-based pricing algorithms also raises significant antitrust concerns. Algorithmic systems can respond to rivals in ways that mimic tacit coordination when repeated interactions generate feedback loops in pricing responses. Airlines that adopt similar AI solutions from shared vendors may also unintentionally coordinate pricing through common algorithms and data inputs. These concerns call for a comprehensive analysis of the product market impacts of AI technologies and their implications for consumer welfare.

[1]The AI systems considered in this paper refer to adaptive, learning-based tools integrated into airlines' pricing and operational processes, rather than the static, rule-based, and domain-limited systems that characterized earlier applications of AI.

[2]Source: Jess Weatherbed, "Delta Air Lines is using AI to set the maximum price we're willing to pay," *The Verge*, July 18, 2025, `https://www.theverge.com/news/709556/delta-air-lines-ai-ticket-price-rollout`.

Our study bridges this gap by examining the impact of AI adoption on product quality, environmental externalities, and pricing in the U.S. airline industry. Specifically, we investigate (1) the effect of AI adoption on quality provisions in terms of on-time performance (OTP) and flight cancellations; (2) its environmental impact, measured by the Air Quality Index (AQI); (3) its effect on pricing behavior and the potential for price coordination; and (4) the resulting implications for market power and consumer welfare.

We construct a panel of U.S. domestic airline operations from 2013Q1 to 2025Q2 and use public financial filings, news reports, and vendor announcements to identify when major carriers adopted AI systems for pricing, scheduling, and maintenance. Leveraging a difference-in-differences (DID) design that compares airline-route pairs before and after the adoption of AI technologies with those that did not adopt, we find that enterprise AI adoption reduces flight delays by an average of 0.4 minutes, is accompanied by a modest increase in cancellations, and improves local air quality by 2.16 AQI units (a roughly 10% improvement). AI adoption also raises average airfares by approximately 1%. The effect reflects a larger increase at the bottom of the fare distribution than at the top, which compresses price dispersion and creates greater uniformity across routes.

To explore potential price coordination channels, we focus on the use of GenAI-enhanced pricing platforms that rely primarily on predictive optimization and reinforcement learning, with generative components serving as supporting tools rather than core decision engines. We then exploit variation in vendor penetration across airlines and routes to examine how coordination arises when carriers rely on a common algorithm vendor and as these platforms spread across airlines. We find that the pricing effects become even larger for GenAI-enhanced pricing platforms, with fares rising broadly across the distribution. However, the direction and magnitude of these effects depend on whether rivals on a given route share the same algorithm vendor and the extent of platform diffusion in the market. In particular, when rivals rely on the same algorithm vendor, fares tend to be higher throughout the distribution, especially at the lower end, whereas the diffusion of GenAI-enhanced pricing

platforms is associated with lower fares across the distribution, with the effect again being more pronounced at the lower end.

To help interpret these opposing effects, we develop a simple theoretical framework in which airlines compete in prices and GenAI platform use affects equilibrium fares through common vendor overlap and overall platform diffusion. In this setting, greater same-vendor overlap dampens competition by expanding the set of rivals whose profits are partially internalized, which raises equilibrium prices. At the same time, the model predicts that as platform diffusion increases within a route, the payoff to further price increases declines, which puts downward pressure on equilibrium fares. These interpretations are also consistent with evidence that greater overlap in common outsourcing relationships can soften competition and raise fares (Aryal et al., 2024), as well as recent work on algorithmic competition and pricing (Musolff, 2025).

Our findings are robust to a range of sensitivity and identification checks that address several threats to interpretation in both the enterprise AI and GenAI analyses. In particular, because AI adoption is observed at the airline-time level while outcomes are measured at the airline-route-time level, the baseline estimates may be affected by staggered treatment timing, alternative ways of measuring AI adoption intensity, and unobserved carrier-level or route-specific shocks that evolve over time and coincide with adoption. For GenAI platform use, an additional concern is that route-level deployment is not directly observed, so airline-level adoption should be interpreted as an intention-to-treat measure rather than a literal route-level treatment. We mitigate these concerns using staggered-adoption estimators, alternative adoption measures, richer fixed-effects and interactive-factor specifications, and falsification exercises that test placebo assignments and whether the results are driven disproportionately by any single airline. Across these exercises, our main results remain broadly consistent.

Our paper makes several contributions to the literature. First, there is a fast-growing body of empirical work examining the economic impact of AI pricing algorithms across

different industries, including retail gasoline (e.g., Assad et al. 2024), hotels and short-term rentals (e.g., Zhang et al. 2021; Chávez-Miranda et al. 2025), e-commerce (e.g., Wang et al. 2024b; Zhao and Berman 2025), and creative goods markets (Goldberg and Lam, 2025). Whereas most existing studies tend to focus on pricing outcomes, our work offers a broader perspective by analyzing airlines' adoption of a wider set of AI tools and their economic impacts, ranging from on-time performance and air quality to fares and their distributions. Our findings in turn provide comprehensive causal evidence of the key trade-off between efficiency gains and coordination risks in the adoption of AI technologies.

To our knowledge, this paper is also among the first to examine the economic consequences of exposure to *generative* AI technologies and distinguish these effects from those associated with enterprise AI. Azoulay et al. (2025) argue that competitive advantage in GenAI depends on control of complementary assets, such as compute and inference infrastructure, safety and governance pipelines, benchmarks, and privileged data with feedback loops, which gives platform vendors scope to shape market structure. Goldberg and Lam (2025) find that GenAI goods substitute for non-GenAI content and crowd out its production, but at the same time improve product quality, expand variety, and increase overall output and sales. Our analysis provides novel empirical evidence from the airline industry that the use of pricing platforms with generative components, introduced as an additional layer on top of existing enterprise pricing systems, can raise fares in settings where rivals share the same vendor, whereas broader diffusion across a market lowers fares. As a result, the welfare impact of GenAI-enhanced pricing platforms depends on how vendor overlap and platform diffusion affect pricing incentives across rivals. These insights also shed light on the mechanisms through which AI adoption may foster coordinated pricing. For example, Kühn and Tadelis (2018) argue that coordination among pricing algorithms is not automatic and instead depends on the informational environment in which firms interact. Our findings are consistent with this view by showing that the pricing effects of GenAI depend on how these tools diffuse across rivals, rather than on their use alone.

Our study complements the diagnostic framework of Harrington (2025b), which provides a test for coordinated adoption of third-party pricing algorithms. Instead of testing for the presence of coordinated adoption itself, we take vendor presence as given and empirically examine how vendor overlap and the diffusion of these pricing tools across carriers and routes influence competitive outcomes. Our results also connect to Hortaçsu et al. (2024), who show that U.S. airline pricing in the late 2010s relied on rule-based heuristics shaped by organizational frictions. In contrast, we analyze the subsequent wave of enterprise and generative AI adoption and document how vendor-driven pricing systems influence pricing patterns and potential coordination channels.

Lastly, our findings provide important welfare and policy implications. A welfare calculation based on our demand estimates indicates that enterprise AI adoption generates welfare gains from improved service quality of about $42.9 million, losses from higher fares of about $3.1 million, and environmental benefits of roughly $0.7 million, yielding a substantial net positive effect for consumers. These estimates suggest that, on balance, AI adoption enhances overall consumer welfare by improving product quality and reducing negative environmental externalities despite modest price increases. From a policy standpoint, however, our evidence also underscores the need for careful regulatory oversight. The risks of algorithmic coordination appear to depend less on the specific type of AI pricing tool than on how it is deployed and the competitive environment. As AI-driven pricing becomes more widespread, regulators should pay close attention to vendor concentration, market structure, and emerging concerns over personalized airfares.[3]

[3]For example, in July 2025 a group of U.S. senators led by Ruben Gallego sent a letter to Delta Air Lines raising concerns that its planned AI-driven pricing could individualize fares based on personal data and push prices to consumers' "pain point." See Bryan Metzger, "Senators demand answers about Delta's new AI-driven pricing plan," *Business Insider*, July 22, 2025, `https://www.businessinsider.com/delta-airlines-ai-pricing-ruben-gallego-letter-2025-7`.

2 Background and Conceptual Framework

2.1 AI Adoption in the U.S. Airline Industry

The U.S. airline industry pioneered forecast-based revenue management systems in the 1980s. By the early 2000s, airlines had begun adopting basic machine learning (ML) algorithms, such as support vector machines and random forests, to improve demand forecasting and fare optimization (Ge et al., 2025b). Since the mid-2010s, they have steadily expanded these efforts through enterprise AI systems across multiple areas of operation. Compared to standalone ML models designed for specific prediction tasks, enterprise AI integrates these algorithms into broader organizational functions and leverages structured data such as operational records, booking histories, demand forecasts, and competitor pricing to inform both pricing and operational decisions. As a result, enterprise AI operates through forecasting cycles and periodic reoptimization and is embedded directly within each airline's proprietary systems.

Delta was among the first airlines to apply enterprise AI in human resource management; in 2014 it replaced its legacy HR infrastructure with SAP Business AI designed to enhance the employee experience and support career development.[4] Another early area of AI adoption is predictive maintenance, which enables carriers to anticipate when aircraft parts require service, reduce downtime, and improve operational efficiency (Ge et al., 2025b). In the following years, AI adoption expanded to flight operations and customer service (see Geske et al. (2024) and Tafur et al. (2025) for a literature review). AI technologies have supported more efficient route planning, scheduling, crew assignments, and baggage tracking, with airlines investing heavily in systems designed to reduce fuel consumption and delays. Airlines also deployed AI chatbots to handle common passenger inquiries and introduced machine learning models to deliver personalized marketing offers and optimize seat selection. More recently, airlines have expanded the use of enterprise AI to pricing and inventory management

[4]Stephanie Ball, "Delta Air Lines Takes Flight with AI Tools," *ERP Today*, n.d., `https://erp.today/delta-air-lines-takes-flight-with-ai-tools/`.

by partnering with algorithm platforms such as PROS, Fetcherr, and FLYR. These systems rely on historical fare data, demand forecasts, and real-time inputs like booking trends and competitor prices to dynamically adjust fares and manage seat inventory (Ge et al., 2025b).

Generative AI has also recently been introduced into airline pricing platforms. Compared to enterprise AI systems that automate pricing and other business functions using predictive modeling, optimization, or reinforcement learning, generative AI produces new content, such as text, code, scenarios, or explanations, based on large language models (LLMs) trained on broad corpora, and can interact with unstructured data (e.g., news, consumer sentiment) to provide summaries and recommendations. In other words, while enterprise AI optimizes and executes pricing and operational decisions, generative AI interprets, communicates, and coordinates those decisions.

In practice, pricing algorithm vendors have incorporated generative AI into their existing enterprise AI platforms as a supporting module rather than a replacement for the core optimization architecture. The diffusion of GenAI technology in the U.S. airline industry thus does not represent a regime change, but rather an incremental layer on top of existing enterprise systems that reduces human-in-the-loop latency, accelerates internal decision cycles, and improves sensemaking across the organization.

It is also worth distinguishing a GenAI-enhanced pricing platform from a fully agentic AI pricing system, which would serve as the primary price-setting mechanism rather than as a workflow layer.[5] Such systems are not yet in use in airline pricing at the time of this writing. Delta's recent initiative, however, points toward this paradigm. While not agentic, Delta launched a generative AI pricing pilot in late 2023 to reduce manual workload and assess the system's ability to identify optimal price points. By mid-2025, Delta reported pricing around 3 percent of its fares using this generative system, with plans to expand coverage to

[5]Such a system would integrate large language model methods with reinforcement learning and neural networks to (1) generate and assess synthetic market scenarios, (2) process structured and unstructured data sources (e.g., demand forecasts, competitor behavior, news, and consumer sentiment), and (3) autonomously set and justify pricing decisions. See Wang et al. (2024a) for an overview of agentic AI systems.

about 20 percent by year-end.[6]

Figure A1 in the Online Appendix presents the adoption timelines of various enterprise AI and GenAI-enhanced systems in the U.S. airline industry. Notably, AI adoption in the airline industry follows a staggered rollout across carriers and routes. Our empirical design exploits this staggered timing and heterogeneity in adoption patterns to identify the product market effects of AI adoption.

2.2 Efficiency and Coordination Effects of AI Adoption

The adoption of enterprise AI systems in the airline industry is expected to improve operational efficiency and reduce the environmental footprint. By optimizing gate assignments, flight scheduling, and route planning, AI technologies can reduce delays and improve on-time performance.[7] In addition, more efficient aircraft movements and predictive maintenance systems can lower fuel consumption and carbon emissions, which in turn reduces local air pollutants and improves air quality. These impacts are likely to be most pronounced at major hubs and on high-frequency routes, where the potential for operational gains is greatest.

There are, however, several channels through which enterprise AI adoption can raise coordination risks in airline pricing. First, algorithmic responsiveness can strengthen parallel price movements when pricing and inventory modules closely monitor rivals and quickly adjust to observed fare changes, especially in environments with high price transparency. Second, common platform exposure can align decision rules when competing carriers rely on the same vendor or similarly designed pricing systems. Third, enterprise AI adoption may lower operating costs, but incomplete pass-through or cost-side distortions can still lead to higher fares without any change in competitive conduct.[8]

[6] Christine Boynton, "Delta Air Lines Begins AI Pricing Experiments," *Aviation Week Network*, December 6, 2023, `https://aviationweek.com/air-transport/airlines-lessors/delta-air-lines-begins-ai-pricing-experiments`.

[7] Even though flight cancellations can impose welfare costs on passengers, a profit-maximizing AI system may utilize cancellations to improve operational efficiency. For example, an AI system that anticipates low load factors may preemptively cancel a flight and reallocate passengers to enhance overall capacity utilization.

[8] Although our study examines potential price coordination, adoption of enterprise AI systems may also

More specifically, because airlines collect detailed data on competitors' pricing, enterprise AI systems trained on these inputs can detect and respond to rivals' fare adjustments. This creates feedback loops in which one airline's price adjustment can trigger parallel responses from others. Over time, the iterative process can stabilize prices across carriers and result in pricing patterns that resemble tacit coordination. Experimental evidence supports this prediction: in simulated airline markets, Gu (2023) shows that deep reinforcement learning pricing algorithms can independently learn to coordinate in duopoly settings, achieving near-monopoly profits more readily than traditional revenue management heuristics. More broadly, controlled experiments suggest that learning-based pricing algorithms can converge to supracompetitive pricing without explicit communication (Calvano et al., 2020; Wang et al., 2024b). Field evidence from the German retail gasoline market also indicates that algorithmic pricing can dampen competition and raise prices (Assad et al., 2024). Additionally, while greater price transparency can enhance efficiency, it can also facilitate coordination by making rivals' fare adjustments easier to detect and match, which reinforces feedback loops and parallel pricing behavior.

Next, coordinated pricing can become more likely when airlines share the same algorithm vendors or similarly designed enterprise AI platforms. Standardized models and optimization routines that draw on overlapping data sources can make carriers facing similar competitive environments set nearly identical prices. By reducing the scope for independent price setting, shared algorithmic platforms may unintentionally make tacit coordination more likely (Ge et al., 2025b).[9]

Furthermore, enterprise AI adoption can affect pricing through cost-side channels. Platforms can lower operating costs by automating routine tasks and improving forecasting and

facilitate coordination in capacity and scheduling. For example, AI tools could lead competitors to make similar adjustments in flight frequency or capacity in response to demand shocks, even in the absence of explicit coordination.

[9] A useful historical analogue is the ATPCO case (*U.S. v. Airline Tariff Publishing Co.*, 1992–94), in which the Department of Justice found that airlines used a common fare-publishing platform to signal planned price increases and monitor whether rivals matched. Such a platform reduced uncertainty and made coordinated pricing easier, all without explicit agreements.

resource allocation, which is consistent with evidence that IT-related outsourcing helps reduce non-IT operating costs (Han and Mithas, 2013). However, the literature on cost pass-through suggests that cost reductions do not translate one-for-one into lower fares when pass-through is incomplete, particularly under imperfect competition (Weyl and Fabinger, 2013). More importantly for coordination risks, cost-side mechanisms can themselves push prices upward. In Harrington (2025a), executives coordinate on inflated internal cost benchmarks that pricing managers subsequently use as inputs in their pricing decisions. This leads to higher prices while avoiding explicit coordination on markups. In this sense, algorithmic coordination can operate through the use of inflated internal cost benchmarks in pricing decisions, rather than through the direct coordination of markups.

2.3 GenAI-Enhanced Pricing Platforms and Coordinated Pricing

Besides investigating how enterprise AI adoption may facilitate price coordination, our study exploits the recent rollout of non-agentic GenAI technology to separately assess two of the aforementioned coordination channels, including algorithmic responsiveness associated with platform diffusion and coordination risk arising from shared algorithm vendors.

Firstly, when rival carriers share the same GenAI-enhanced pricing platform, their algorithms are more likely to align, which can narrow the scope for independent price setting and raise coordination risk. This mechanism is analogous to Aryal et al. (2024), who demonstrate that major carriers subcontracting to the same regional airline exhibit higher average airfares. At the same time, using the same algorithm vendor may also lower fares when the vendor internalizes competition among its clients and avoids embedding aggressive high-price algorithms (Harrington, 2022). In addition, when competing carriers rely on similar automated pricing tools that react to competitors' posted fares, rapid undercutting can trigger price wars akin to Edgeworth price cycles, especially early in the diffusion process (Musolff, 2025). Taken as a whole, the net effect of shared GenAI-enhanced platforms is an empirical question depending on how vendor incentives and algorithmic design affect pricing

responses. Compared to earlier pricing systems, a GenAI layer may further magnify these effects by reducing internal decision lags and accelerating how pricing signals are interpreted, communicated, and acted upon within firms.

Second, coordination risk can also arise as GenAI-enhanced pricing platforms diffuse more broadly across carriers and routes, even when rivals rely on different vendors. In a seminal study, Calvano et al. (2020) run controlled experiments with Q-learning algorithms and show that even relatively simple reinforcement learning algorithms can converge to supra-competitive prices without explicit communication. While GenAI is not itself a learning algorithm, the experimental evidence in Calvano et al. (2020) provides useful intuition for how faster and more automated pricing responses can stabilize pricing outcomes. In particular, as GenAI-enhanced pricing platforms diffuse across carriers on a route, pricing becomes more systematic and predictable. Because vendor systems are built on similar optimization principles, they tend to respond in parallel to common market signals, which reduces uncertainty about rivals' pricing behavior and makes aggressive undercutting less likely (Wang et al., 2024b). Such stability can facilitate parallel pricing behavior and lead to lower fare dispersion. Meanwhile, widespread diffusion of algorithmic pricing platforms can generate operational efficiencies by improving the speed and accuracy of demand forecasting and price adjustments across the market. Such efficiency gains can lower average fares by allowing airlines to better align prices with demand. As with the shared vendor channel, the net price effect of algorithmic responsiveness is therefore an empirical question. However, unlike the shared algorithm vendor channel, which operates through alignment among rivals using the same pricing system, the platform diffusion channel reflects broader market-wide effects that arise as automated pricing becomes more widespread across carriers and routes.

3 Data

3.1 Data Sources and Sample Construction

We construct a panel of U.S. domestic airline operations from 2013Q1 to 2025Q2 at the airline–route–month level. The timelines of airline AI adoption are identified using information from public financial filings, news reports, and vendor announcements that document when major U.S. carriers implemented various AI and GenAI-enhanced systems for pricing, scheduling, and maintenance.[10]

Our primary measures for an airline's quality provisions focus on its on-time performance record, which can be directly affected by the carrier's route-specific investments (Mazzeo 2003; Prince and Simon 2015). Data on airline on-time performance are obtained from the On-Time Performance database provided by the Bureau of Transportation Statistics (BTS). The OTP data contain daily flight-level observations and rich information about arrival and departure delays, types of delays, and flight cancellation rates. Following related literature, we define the product market as directional airport pairs and exclude routes from and to Alaska, Hawaii, and Puerto Rico.

Our main OTP measures are *Delay Time*, defined as the total delay minutes per flight after excluding weather-related disruptions, and *Cancellation*, defined as the share of flights canceled for reasons other than weather. By excluding weather-related delays and cancellations, these measures capture operational disruptions driven by airline performance rather than exogenous weather shocks.

Airfare data are obtained from the BTS's DB1B database, which represents a quarterly 10% random sample of all domestic airline tickets sold in the U.S. In addition to measuring average airfares, we follow prior literature on price dispersion in the airline industry (e.g., Borenstein and Rose 1994; Gerardi and Shapiro 2009; Kim et al. 2023) and include measures of a carrier's fare distribution on a given route based on the 20th and 80th percentiles of the

[10] See Online Appendix A for examples of adoption timelines and sources.

price distribution, which are respectively used to capture leisure vs. business travelers.

Air Quality Index data come from the Air Quality System (AQS) maintained by US Environmental Protection Agency (2025). The EPA's air quality monitors report daily AQI values along with the GPS coordinates of each monitoring location. We match each airport to its nearest monitoring station and include the quarterly average AQI for both origin and destination airports.

3.2 Vendor Overlap and Platform Diffusion Measures

We construct two route–time–level indices to capture the extent of vendor overlap and the degree of diffusion of GenAI-enhanced pricing platforms across airlines on a given route. In particular, the Shared-Vendor Overlap Index (SVOI) measures the extent to which competing airlines on a route use the same AI vendor and captures the potential for coordinated pricing arising from the use of a common third-party provider such as PROS or Fetcherr.[11] The Vendor Adoption Intensity Index (VAII), on the other hand, captures the intensity of third-party pricing vendor usage, or how algorithmic the pricing is on a given route, regardless of the specific vendor.[12]

More formally, in our carrier-route panel, each airline i serving a directional route j at time (year-quarter) t may rely on a third-party algorithm vendor $v_{it} \in \mathcal{V} \cup \{\text{None}\}$, where $\mathcal{V}$ is the set of vendors including PROS and Fetcherr and None indicates that the carrier relies on an in-house or legacy system that is not shared with rivals. Further denote

[11] Note that algorithm vendors such as PROS and Fetcherr license their pricing platforms to airlines on an individual-client basis, and public documentation describes airline-specific implementations. We are not aware of any publicly documented case in which proprietary pricing or booking data from one airline were used to set prices for another. Airlines also emphasize data separation, e.g., Delta has publicly stated that it does not share personal information with Fetcherr and its pricing system does not use personal data (Delta Air Lines, "Delta responds to misinformation around AI pricing," *Delta News Hub*, August 7, 2025, `https://news.delta.com/delta-responds-misinformation-around-ai-pricing`). While vendors may update models using aggregate performance insights across clients, how or whether learning transfers across airlines is not disclosed. Accordingly, our SVOI measure captures shared-vendor status (common algorithm software and vendor-updates), not cross-airline data sharing.

[12] Also note that conversations with operations and engineering personnel at a major U.S. carrier indicate that in-house AI pricing tools are actively under development but not yet deployed at scale. To the extent that such systems remain experimental or limited in scope during our sample period, the SVOI and VAII measures effectively capture the primary channels through which GenAI enters airline pricing.

$\text{Alg}_{it} \equiv \mathbb{1}\{v_{it} \neq \text{NONE}\}$ as an indicator for whether airline i employs any third-party vendor at time t, and $s_{ijt} \in [0,1]$ as the market share of airline i on route j at time t, measured in terms of enplanements and normalized so that $\sum_{i \in \mathcal{K}_{jt}} s_{ijt} = 1$.

Let $w_{ii't} \geq 0$ be a pairwise weight for airlines $i \neq i'$ serving route j, and $w_{ii't}$ is set equal to the fraction of routes that airlines i and i' operate in common relative to the total number of distinct routes operated by the two carriers in a given period. Thus, pairs of airlines with greater network overlap receive larger weights. We define the Shared-Vendor Overlap Index as the weighted average of an indicator for whether two airlines use the same algorithm vendor:

$$\text{SVOI}_{ijt} = \frac{1}{W_{ijt}} \sum_{v} \sum_{i \neq i', i, i' \in J_{jt}} w_{ii't} \times \mathbb{1}\big\{ v_{it} = v_{i't} \big\}, \tag{1}$$

where $\mathcal{J}_{jt}$ is set of airlines serving route j at time t and $W_{ijt} = \sum_{i \neq i', i, i' \in J_{jt}} w_{ii't}$. By construction, $\text{SVOI}_{ijt} \in [0,1]$, and the measure captures the fraction of (weighted) airline pairs serving route j at time t that rely on the same algorithm vendor. Thus, a higher value of SVOI indicates greater alignment on the same vendor. Note that if fewer than two airlines serve route j at time t, we set $\text{SVOI}_{ijt} \equiv 0$.

In our setting, United and JetBlue are the only major carriers that share PROS as a common AI pricing vendor, while Delta relies on Fetcherr and other carriers employ in-house systems. The SVOI measure can therefore be written as

$$\text{SVOI}_{ijt} = \frac{1}{W_{ijt}} \times w_{UA,B6,t} \times \mathbb{1}\big\{ \text{UA, B6} \in J_{jt} \big\}.$$

Analogously, we define the Vendor Adoption Intensity Index as the share of airlines on route j at time t that use any third-party AI pricing vendor (regardless of which one):

$$\text{VAII}_{jt} = \frac{1}{|J_{jt}|} \sum_{i \in J_{jt}} \mathbb{1}\big\{ v_{it} \neq \text{NONE} \big\}. \tag{2}$$

By construction, $\text{VAII}_{jt} \in [0,1]$, and the measure captures the fraction of airlines on route j

at time t that rely on a third-party algorithm vendor. VAII therefore reflects the degree of algorithmic intensity, or "platformization," of pricing in the market. Note that if no vendors serve route j at time t, we set VAII to 0 by convention.

Conceptually, VAII is distinct from SVOI, e.g., adding a rival that uses a different AI pricing vendor increases VAII but leaves SVOI unchanged. While SVOI captures the potential coordination risk due to shared pricing platforms, VAII measures how widespread AI pricing has become on a route and how pricing outcomes evolve as a larger share of airlines rely on GenAI-enhanced pricing platforms. Online Appendix B provides further examples and computational steps for the SVOI and VAII measures. Section 6.3 discusses potential endogeneity concerns related to these indices.

3.3 Descriptive Patterns

Columns 1–3 of Table 1 report the summary statistics for the dataset used for the enterprise AI adoption regressions. We restrict the sample through 2023Q3 to exclude the period when airlines started adopting GenAI. Here, we observe that average fares are modestly higher on adopted routes than on non-adopted routes, and the gap appears more pronounced at the upper end of the fare distribution, resulting in greater price dispersion on adopted routes. Product quality indicators also differ by AI adoption status. Total delay time is noticeably shorter on adopted routes, whereas flight cancellation rates are higher. This pattern is consistent with operational reoptimization, where airlines prioritize reducing extended delays potentially at the cost of more cancellations. Local air quality, on the other hand, is broadly comparable between the two groups.

Columns 4–6 present the summary statistics for the dataset used for the GenAI platform use regressions. Routes with GenAI platform use have substantially higher average fares than routes without such use, and the difference appears at both the lower and upper ends of the fare distribution. The adopted group also exhibits higher SVOI and VAII values, which imply greater vendor overlap and wider platform diffusion on those routes. At the

same time, the overall levels of SVOI and VAII in the full sample remain low, which is consistent with the early stage of GenAI rollout and the limited extent of vendor overlap and platform diffusion during our sample period.

Figure 1 presents the breakdown of service quality and fares by enterprise AI adoption status and GenAI platform use measures. Panel (a) shows that AI-adopted routes have noticeably better on-time performance and higher cancellation rates. Panel (b) indicates somewhat higher average fares on adopted routes, and the gap becomes larger at the upper end of the distribution. Panel (c) plots fares against SVOI and shows substantial dispersion with no clear systematic relationship between vendor overlap and prices. Likewise, Panel (d) plots fares against VAII and reveals no strong association between the diffusion of GenAI-enhanced pricing platforms and fare levels.[13]

4 Empirical Specifications

To identify the causal impact of AI adoption, we employ a difference-in-differences design that compares airline–route pairs that adopted enterprise AI with those that did not across the pre- and post-adoption periods. Specifically, we estimate the following specification:

$$Y_{ijt} = \alpha + \beta \cdot \text{AI}_{ijt} \times \text{Post}_t + \gamma X_{ijt} + \mu_i + \lambda_j + \delta_t + \varepsilon_{ijt}, \tag{3}$$

where Y_{ijt} is the dependent variable of interest for airline i on route j at time t, including on-time performance, cancellation rate, AQI, log average fare, and the 20th and 80th percentile fares. AI_{ijt} is an indicator for whether airline i has adopted enterprise AI tools on route j by time t, and $Post_t$ is a post-treatment time dummy. X_{ijt} includes time-varying controls such as airport network size, route concentration (HHI), and route distance.

Our baseline specifications include airline fixed effects μ_i, route fixed effects λ_j, and

[13] Additionally, Figure A2 in the Online Appendix plots the evolution of airfares, market concentration, OTP, and AQI, while Appendix Figure A3 shows the distribution of SVOI and VAII.

time (year-quarter) fixed effects δ_t to control for airline-specific factors, time-invariant route characteristics, and common temporal shocks, respectively. Standard errors are clustered at the route level to account for serial correlation and policy exposure. Note that we are unable to include carrier–time fixed effects because they are perfectly collinear with the treatment variable. Nevertheless, we conduct several related robustness checks to address this concern. First, we interact route fixed effects with linear time trends, and separately interact carrier fixed effects with linear time trends, to account for differential trends across routes and airlines that could otherwise confound the estimates. Second, we re-estimate the main specifications using the interactive fixed-effects framework by Bai (2009), which allows latent common shocks to affect observational units differently over time and therefore provides a more flexible way to capture time-varying unobserved heterogeneity beyond additive fixed effects. These robustness checks are discussed in more detail in Sections 6.4 and 6.5.

The parameters in Equation 3 are identified from within–airline-route variation over time, conditional on observed route characteristics X_{ijt} and airline, route, and time fixed effects. A key identifying assumption is that in the absence of enterprise AI adoption, treated and untreated airline–routes would have followed parallel trends in fares, service quality, and environmental outcomes. This common trends assumption is plausible in our empirical context because the timing of adoption is largely driven by firm-level technology partnerships, vendor rollout schedules, and broader industry-wide trends toward AI adoption, rather than by short-term demand, weather, or other route-specific shocks.[14] In Section 6.1, we further assess this assumption using a formal event-study analysis with leads and lags of adoption.

To account for potential non-random adoption of enterprise AI, we adopt a reduced-form two-step approach motivated by Sweeting et al. (2020) that estimates a first-stage probit model and incorporates the fitted adoption probability into the baseline specifica-

[14]Note that the enterprise AI deployments we study were not accompanied by major reorganizations, bankruptcies, or broad operational overhauls. Conversations with operations and engineering personnel at a major U.S. carrier also indicate that AI was implemented as a standalone technology initiative rather than part of a broader reshuffling of investments or strategy. This mitigates concerns that our estimated effects are driven by contemporaneous carrier-wide changes unrelated to AI adoption.

tion. Because airlines adopt enterprise AI systems across their operations rather than on a route-by-route basis (e.g., firm-wide customer-facing AI tools or pricing systems), adoption occurs at the airline-time level. Accordingly, we estimate the first stage using airline-time observations rather than airline–route–time observations. The dependent variable is an indicator that equals 1 if airline i has adopted enterprise AI by time t and 0 otherwise, and the first-stage regressors consist of lagged airline characteristics, including fare, administrative expense, reservation expense, line service expense, number of routes, average distance, and total tickets sold. We then include the fitted adoption probability, $\widehat{\Pr(\text{AI adoption})}_{it}$, as a generalized regressor in Equation 3 to capture systematic differences in airlines' underlying propensity to adopt enterprise AI. Because adoption is determined at the airline level, this estimated propensity reflects observable carrier characteristics associated with adoption that may also influence route-level outcomes. Conditioning on this propensity therefore helps account for selection on observable airline characteristics that may otherwise confound the estimated effects of AI adoption on route-level outcomes.

To investigate potential price coordination channels, we focus on carriers that rely on GenAI-enhanced pricing platforms and comparable airlines that do not and augment Equation 3 by incorporating SVOI and VAII indices. Specifically, we estimate

$$\begin{aligned} P_{ijt} = \alpha + \beta_1 \text{GenAI}_{it} + \beta_2 \text{SVOI}_{ijt} + \beta_3 \text{VAII}_{jt} + \beta_4 (\text{GenAI}_{it} \times \text{SVOI}_{ijt}) \\ + +\beta_5 (\text{GenAI}_{it} \times \text{VAII}_{jt}) + X'_{ijt}\gamma + \mu_i + \lambda_j + \delta_t + \varepsilon_{ijt}, \end{aligned} \tag{4}$$

where P_{ijt} denotes average fares or fare percentiles, and $GenAI_{it}$ equals 1 for airline i that uses a GenAI-enhanced third-party pricing platform at time t and thereafter, and 0 otherwise. Because the route-level implementation of GenAI technology is not directly observed, we interpret $GenAI_{it}$ as an intention-to-treat (ITT) indicator that captures airlines' adoption of GenAI-enhanced pricing infrastructure, rather than route-level or flight-level implementation of GenAI pricing.

The coefficients in Equation 4 capture how GenAI platform use interacts with vendor commonality and platform diffusion. β_1 measures the baseline effect of relying on a GenAI-enhanced pricing platform in markets with no AI pricing vendor presence. β_2 represents the effect of shared-vendor overlap when the airline has not yet transitioned to a GenAI-enhanced pricing platform. β_3 captures the effect of overall vendor diffusion in the absence of GenAI platform use. The interaction terms measure additional effects associated with GenAI-enhanced platform use. β_4 captures the additional effect of GenAI platform use in markets with high vendor overlap, while β_5 reflects the additional effect of GenAI platform use in markets with high overall vendor diffusion.

One may be concerned that the limited number of GenAI pricing platform adopters may constrain identification. We, however, argue that the concentrated pattern of vendor adoption in fact helps provide relatively clean variation for identifying price coordination channels. In particular, vendor choice in the airline industry is determined at the carrier level based on long-term technology partnerships and implementation capacity, rather than route-level competition. United and JetBlue's shared adoption of PROS and Delta's separate adoption of Fetcherr therefore reflect carrier-level decisions that are plausibly exogenous to the route-level pricing outcomes we study. This differs from alternative empirical settings in which several AI pricing vendors compete and vendor choice becomes endogenous to market structure. Additionally, the limited extent of GenAI diffusion in our sample reflects the early-stage rollout of these technologies in a capital-intensive industry with legacy pricing systems. Studying this early diffusion stage allows us to identify the incremental impact of GenAI platform adoption before widespread adoption moves the market toward a new pricing equilibrium.

5 Impacts of AI Adoption and Diffusion

5.1 On-Time Performance

In Table 2, columns 1–2 present the estimates for flight delay time and cancellation rates after excluding weather-related disruptions. We observe that enterprise AI adoption decreases average delay time by 0.39 minutes and increases the likelihood of flight cancellation by 7 percentage points. This result likely reflects operational reoptimization, where airlines adopting enterprise AI prioritize minimizing extended delays and use cancellations to rebalance operations. Columns 4–5 report estimates from the generalized regressor approach that incorporates the predicted adoption probability as outlined in Section 4, and the results are almost identical to the OLS estimates.[15]

It is worth noting that in a hub-and-spoke network, operational spillovers across connected routes imply that improvements (or disruptions) propagate beyond the treated route, which can attenuate route-level estimates. To the extent that these spillovers operate in the same direction, our route-level estimates should therefore be interpreted as a lower bound on the true effects of enterprise AI on operational outcomes.

5.2 Local Air Quality

Columns 3 and 6 of Table 2 further suggest that enterprise AI adoption can lead to significant improvements in local air quality, with estimated reductions in AQI of about 2.2 points in both the OLS and generalized regressor specifications. Since AQI is a scaled index where lower values indicate cleaner air, these estimates imply an economically significant improvement of roughly 10 percent relative to the sample mean of 22.55 shown in Table 1. In other words, airlines' enterprise AI adoption is associated with not only improved service quality through better on-time performance but also measurable reductions in air pollution

[15]The first-stage results are reported in Appendix Table A1.

exposure for communities surrounding the airports.[16]

5.3 Average Fares and Price Distribution

Our DID estimates in Table 3 indicate that enterprise AI adoption increases average airfares by 1%. This increase is consistent with potential coordination in airline pricing (Aryal et al., 2024). The effect is more pronounced at the lower end of the fare distribution, where prices increase by 2.5%, compared to just 0.8% at the higher end. This implies that AI adoption reduces the dispersion of airfares across routes by raising lower-end prices while keeping higher fares relatively stable. The generalized regressor estimates from columns 4–6 closely corroborate the OLS results and exhibit nearly identical patterns for average fares and fare distribution. These results thus suggest that enterprise AI adoption is associated with both an overall increase in airfares and a compression of the fare distribution, patterns consistent with potential price coordination.

5.4 GenAI Platform Use, Common Vendor, and Pricing

Table 4 examines the price effects of GenAI platform use. Columns 1–3 present the baseline ITT estimates and indicate that GenAI platform use on a given route raises average fares by about 5%. The effect is comparable at both ends of the fare distribution. This suggests that GenAI platform use raises prices across the fare distribution without compressing the distribution. Compared to Table 3, the estimates in Table 4 indicate considerably larger and more uniform price impacts across the board.

Columns 4–6 present the estimates from Equation 4 that incorporates the vendor commonality and platform diffusion measures. Focusing on column 4, the main effects of SVOI and VAII are 0.430 and -0.156, respectively, indicating that greater common-vendor overlap

[16]Similar to Azar et al. (2018), our analysis does not explicitly model carriers' capacity choices. A potential concern is that improvements in on-time performance or local air quality could reflect reduced flight activity rather than operational efficiency. To address this, we regress the number of departures and total seats on AI adoption. As reported in Table 11 and discussed in Section 7.1, capacity in fact increases following enterprise AI adoption, indicating that the observed improvements are unlikely to be driven by reduced flights.

(platform diffusion) is associated with higher (lower) average fares. The coefficient on *GenAI* indicates that when vendor overlap is absent and platform diffusion is minimal (i.e., when SVOI = VAII = 0), GenAI platform use raises average fares by approximately 7.7%. This effect increases with common-vendor overlaps (mainly United-JetBlue on PROS), as shown by the positive coefficient on the $GenAI \times SVOI$ interaction term (0.209), but the coefficient is not precisely estimated. In contrast, platform diffusion also reduces the incremental impact of GenAI platform use, as reflected in the negative coefficient on the $GenAI \times VAII$ interaction term (-0.025).

Column 4 further reports the marginal effects of *GenAI*, *SVOI*, and *VAII*, and we find them broadly consistent with the baseline effects. In particular, because SVOI ranges between 0 and 0.14 as shown in Appendix Figure A3, its estimated marginal effect of 0.45 implies an economically meaningful effect as a 0.05 increase in SVOI, roughly the difference between markets with minimal and moderate same-vendor overlap, is associated with about a 2.3% increase in fares. The marginal effect of VAII remains negative, which suggests that greater platform diffusion lowers average fares while also reducing the incremental effect of GenAI platform use. A 0.1 increase in VAII, which ranges between 0.1 and 1.0 across route–time observations, corresponds to a 1.58% decrease in average fares. This implies that moving from low to high platform diffusion (0.3 to 0.9) is associated with an economically significant 9% fare decline.

Columns 5–6 extend the results to the fare distribution. We observe that GenAI platform use raises high-end fares more than low-end fares, with the marginal effects indicating increases of about 6.4% at the upper end and 4.3% at the lower end. Common-vendor overlap dampens this effect, especially at the lower end of the fare distribution, as reflected in the negative $GenAI \times SVOI$ coefficients in both columns and the much larger negative interaction for low-end fares. Greater platform diffusion is associated with lower fares across the distribution, and the effect is again more pronounced at the lower end. Its interaction with GenAI, however, differs across the distribution, with a positive $GenAI \times VAII$ coefficient

for low-end fares but a negative one for high-end fares. The marginal effects support this interpretation: GenAI platform use raises fares throughout the distribution, SVOI is associated with higher fares, especially at the lower end, and VAII is linked to lower fares across the distribution, with the decline again being larger at the lower end.

Taken as a whole, the results from Table 4 indicate that GenAI platform use can raise prices in isolation, but its ability to sustain coordinated pricing depends on vendor choice and the extent of platform diffusion in the market. The estimates point to two distinct effects. Consistent with a coordination-based interpretation, greater reliance on a common vendor among competitors is associated with higher fares, particularly through its effect on the lower end of the fare distribution. In contrast, markets with more widespread platform diffusion tend to exhibit lower fares, especially at the lower end of the fare distribution.

6 Robustness Checks

In this section, we discuss a range of robustness checks to assess the sensitivity of our main results to alternative specifications and sample restrictions.[17]

6.1 Staggered AI Adoption

While our baseline DID specification employs carrier, route, and time fixed effects, the staggered timing of AI adoption across airlines and routes implies that time fixed effects may be estimated using both treated and untreated units at each point in time. This can bias two-way fixed effects (TWFE) estimates when treatment effects vary across cohorts or over time. To address this concern, we follow recent advances in DID with staggered adoption (Goodman-Bacon, 2021; Sun and Abraham, 2021; Callaway and Sant'Anna, 2021) and implement the Callaway and Sant'Anna (2021) estimator. The CS estimator identifies group–time

[17]Note that because the enterprise AI and GenAI analyses involve somewhat different measurement and identification concerns, some robustness checks apply to both sets of results, while others are specific to one setting.

average treatment effects for each cohort and time period by comparing treated observations to appropriate control groups consisting of not-yet-treated and never-treated units. These group–time effects are then aggregated to obtain overall and dynamic treatment effect estimates. This method yields estimates that are robust to treatment effect heterogeneity and avoid contamination from already-treated units being used as controls. Figure 2a plots the dynamic treatment effects of enterprise AI adoption on airfares relative to adoption timing.[18] We observe that the pre-adoption dummies are statistically insignificant, which provides evidence consistent with parallel pre-trends and the absence of anticipatory pricing behavior prior to adoption.

6.2 Continuous AI Adoption Measures

As a robustness check, we replace the discrete treatment indicator in the baseline specifications with a continuous measure based on the share of airlines on each route that have adopted enterprise AI and re-estimate the corresponding route-level pricing regressions. Since the enterprise AI adoption data are at the airline–time level, the results in Table 5 allow us to examine the impact of route-level AI adoption intensity on airline pricing. Column 1 indicates that a higher AI adoption share is associated with lower average fares, while the positive interaction between enterprise AI share and lagged HHI suggests that this effect becomes less negative and can turn positive on more concentrated routes. Columns 2–3 show similar patterns for lower-end and higher-end fares. Although the interaction effect is somewhat larger for lower-end fares, the estimates are broadly similar across the two ends of the fare distribution and indicate only limited effects on fare dispersion.

Columns 4–6 of Table 5 report the corresponding estimates based on GenAI platform use intensity. Consistent with the enterprise AI results, column 4 indicates that a higher GenAI use share is associated with lower average fares, with the positive interaction term implying

[18]The dynamic treatment effect analysis focuses on enterprise AI adoption because GenAI platform adoption occurs relatively late in our sample period, leaving limited post-adoption periods for identifying informative dynamic effects.

that this effect diminishes and may reverse in more concentrated routes. Columns 5–6 show comparable patterns across the fare distribution and again suggest limited effects on fare dispersion.

6.3 Potential Endogeneity of SVOI and VAII

Because only three airlines rely on GenAI-enhanced pricing platforms, with United and JetBlue using the same vendor (PROS) and Delta relying on a different platform (Fetcherr), variation in vendor overlap and platform diffusion is inherently limited. As a result, the estimates primarily capture effects among early adopters rather than a fully diffused equilibrium. Meanwhile, the small number of vendors reduces concerns about endogenous vendor choice, since vendor selection is essentially binary and largely driven by strategic partnerships with the vendor rather than route-level competitive considerations.

Nevertheless, to address the potential endogeneity of the SVOI and VAII measures, we adopt a control function approach in specifications that include these indices. Specifically, we estimate the following first-stage regressions:

$$\ln \text{SVOI}_{ijt} = \beta_0 + \beta_1 \sum_{i' \neq i, i \in \mathcal{J}_t} \ln \text{SVOI}_{i'j,t-1} + \mathbb{X}'\theta_{ijt}$$

$$\ln \text{VAII}_{ijt} = \alpha_0 + \mathbb{X}'\gamma_{ijt},$$

where the set of control variables $\mathbb{X}$ includes lagged administrative expenses, reservation and sales expenses, and line service expenses, each normalized by its yearly mean and standard deviation.

In the second stage, we re-estimate Equation 4 excluding the interaction terms, while retaining *SVOI* and *VAII* in the equation and including the residuals from the *SVOI* and *VAII* first-stage regressions as additional controls. The estimated signs and magnitudes in Table 6 remain consistent with the marginal effects reported in Table 4. Overall, these results indicate that our main findings are robust to accounting for the potential endogeneity

of SVOI and VAII through the control function approach.

6.4 Alternative Fixed Effect Structures

To address the concern that AI adoption may proxy for broader airline modernization efforts or other persistent differences across airlines and routes, we estimate a sequence of specifications increasingly stringent fixed effect structures. In particular, Table 7 shows that the estimated effect of enterprise AI adoption on fares remains significantly positive after controlling for airline fixed effects, route fixed effects, and time fixed effects, together with either route-specific linear trends, airline-specific linear trends, or airline-by-route fixed effects. These specifications tighten the source of identification in different ways. The models with airline-specific and route-specific trends absorb differential changes in pricing at the carrier and route levels, so the remaining identifying variation comes from deviations from those underlying trajectories that coincide with adoption timing rather than gradual changes in airline modernization efforts or pricing strategies. Our most stringent specification includes airline-by-route fixed effects, so the coefficient is identified from within-airline-route changes over time, comparing the same carrier on the same route before and after AI adoption while netting out common aggregate shocks through time fixed effects. As a result, the estimates here do not rely on permanent differences in carrier quality, route composition, network position, pricing sophistication, or other time-invariant airline-route characteristics.

Across all specifications in Table 7, the estimated effects of enterprise AI adoption on average fares and the lower and upper ends of the fare distribution remain similar to our baseline results in Table 3. Although the magnitudes vary somewhat across specifications, the estimates remain positive throughout for average fares as well as for both ends of the fare distribution. These results thus indicate that our main findings are not driven simply by persistent differences across adopting and non-adopting carriers or by gradual carrier- or route-level trends. Although adoption occurs at the airline-time level, airline-by-route fixed effects absorb only time-invariant heterogeneity for each carrier-route pair and do not

capture changes over time in treatment status. Once a carrier adopts AI technologies, all routes served by that carrier switch treatment status in the post-adoption period, so the estimates are identified from within-carrier-route changes relative to pre-adoption outcomes. For the GenAI platform use analysis, we estimate analogous specifications with airline-by-route fixed effects, route fixed effects, and time fixed effects. The results in Table 8 remain consistent with those reported in Table 4, indicating that the GenAI findings are likewise robust to controlling for common time shocks and time-invariant heterogeneity at the carrier-route level.

6.5 Interactive Factor Fixed Effects

Another concern in our empirical setting is that AI adoption is observed at the airline-time level, whereas the main outcome variables are measured at the airline-route-time level. This creates two related challenges. First, because airline-by-time fixed effects would absorb airline-level adoption indicators, our baseline specifications cannot fully control for all airline-specific shocks that evolve over time, such as unobserved managerial decisions about whether, when, and how aggressively to adopt AI technologies. Second, for GenAI-enhanced pricing platforms, route-level deployment is not directly observed. Our main GenAI specification therefore interprets airline-level GenAI adoption as an intention-to-treat measure that captures the adoption of pricing infrastructure rather than a literal route-level treatment. These complications raise the possibility that our baseline estimates could partly reflect unobserved carrier-level managerial, organizational, or technological changes that coincide with AI adoption.

To mitigate this concern, we re-estimate our main fare specifications using the interactive fixed-effects framework of Bai (2009). This approach augments the standard additive fixed-effects model with a latent factor structure that allows unobserved common shocks to affect observational units differently over time. In our setting, this approach is useful because it can flexibly absorb latent time-varying heterogeneity that may otherwise be correlated with

AI adoption but is not captured by airline, route, and time fixed effects alone. Because the estimator is computationally intensive, we implement it on restricted subsamples, limiting the enterprise AI sample to routes with more than 1,000 passengers and the GenAI sample to routes with more than 500 passengers.

Table 9 reports the resulting estimates. For enterprise AI adoption, the coefficients remain significantly positive for average fares as well as both ends of the fare distribution. The estimated effects are 1.8% for average fares, 2.5% for the 20th percentile, and 1.6% for the 80th percentile. These results are qualitatively consistent with our baseline findings and indicate that the positive fare effects of enterprise AI adoption are not driven solely by unobserved time-varying shocks omitted from the baseline additive fixed-effects specification.

For GenAI platform use, the interactive-factor estimates appear smaller than in the baseline specification, but the positive effect on average fares remains statistically significant. In particular, GenAI platform use is associated with a 2.1% increase in average fares and a 0.7% increase in lower-end fares, while the estimate for upper-end fares is close to zero and not precisely estimated. Thus, after allowing for richer latent unobserved heterogeneity, our results still points to higher average fares following GenAI platform adoption, although the distributional effects are more mixed than in Table 4.

Overall, the interactive fixed-effects estimates in Table 9 confirm that the main price results are not specific to the baseline additive fixed-effects specification.

6.6 Leave-One-Airline-Out Regression

Table 10 reports a leave-one-airline-out robustness exercise in which we re-estimate our baseline GenAI specification by excluding one airline at a time from the sample while keeping all other aspects of the specification unchanged.[19] This design allows us to assess whether the estimated relationship between GenAI platform use and airfare distribution is driven

[19]Note that this robustness check focuses on GenAI platform use because it captures a more comparable pricing-related application across airlines, whereas enterprise AI adoption reflects different mixes of AI functions across carriers, making the leave-one-airline-out results harder to interpret.

disproportionately by any single carrier. Across all specifications, the estimated coefficients on *GenAI* remain qualitatively similar to those in the baseline results, and they continue to indicate a positive association between GenAI platform use and average fares, with corresponding effects across the fare distribution that are generally in line with our main findings in Table 4. These results therefore suggest that our baseline GenAI findings do not hinge on the inclusion of any particular airline in the sample.

6.7 Placebo Test on AI Adoption Timing and Additional Robustness Checks

To verify that our empirical specifications do not produce spurious effects in the absence of true treatment, we conduct separate placebo tests for enterprise AI adoption and GenAI platform use by randomly reassigning treatment status across observations. For each placebo test, we repeat the randomization 100 times and re-estimate Equation 3 for each draw. Figures 2b and 2c plot the respective distributions of the estimated coefficients for AI adoption and GenAI platform use, and we find both distributions approximately normal and centered at zero, indicating that our main results are unlikely to be driven by spurious treatment assignment.

Lastly, in Online Appendix C, we also conduct two additional robustness checks addressing potential measurement error in adoption timing and route-level COVID-19 disruptions. We find our main estimates remain robust across these exercises.

7 Discussion

7.1 Interpreting Pricing Effects of AI Adoption

Our empirical findings indicate that, while enterprise AI adoption improves service quality and local air quality, it is also associated with higher average fares and lower price disper-

sion. Several alternative mechanisms could account for this pricing pattern, including (1) quality improvements; (2) capacity reductions; (3) revenue management optimization; and (4) algorithmic coordination. This section evaluates each channel in turn.

7.1.1 Quality Improvements

One potential explanation consistent with the results reported in Table 3 is that enterprise AI adoption improves service quality and operational reliability, which raises passengers' willingness to pay and thus airfares. We outline three empirical tests to distinguish price effects from quality improvements.

First, we regress fares on OTP–related variables, including lagged delay time and cancellation rates, with the full set of controls and fixed effects used in Equation 3. We then use the residuals from this regression, which capture the portion of fares not explained by service quality improvements, as a proxy for markups. Next, we regress these residuals on the AI adoption indicator and the same set of fixed effects and report the results in Table 11. Column 1 exhibits a positive effect of enterprise AI adoption on the markup proxy, indicating that even after accounting for service quality gains, prices still rise with AI adoption. This is also consistent with the price effects operating through a markup or conduct channel and supports the coordination interpretation.

Next, we include lagged flight cancellation rates as an additional control in our baseline price regression. Column 2 of Table 11 indicates that the coefficient on AI adoption stays positive and significant even after accounting for lagged cancellations. This suggests that the remaining price effect likely reflects increased market power rather than higher willingness to pay from improved service quality.

In addition, we compare the residuals from the baseline price regression with those from a specification that excludes enterprise AI adoption. The difference would isolate the portion of fares attributable to AI adoption. We apply the same procedure to delay time and flight cancellations to construct $\Delta Delay$ and $\Delta Cancellation$, which capture changes in service

quality associated with AI adoption. We then test whether routes and carriers with greater improvements in OTP also experience larger fare effects by regressing the fare residual differences on these quality measures using airline-route averages. Column 3 of Table 11 shows that both coefficients are insignificant, which indicates that the fare effects are not driven by improvements in on-time performance. Taken together, the results from columns 1 through 3 of Table 11 suggest that the observed price effects of enterprise AI adoption are unlikely to be driven primarily by higher willingness to pay for improved service quality and are instead more consistent with coordination channels.

7.1.2 Capacity Reductions

A second potential mechanism is that enterprise AI adoption may lead airlines to reduce flights and seat capacity, which could raise equilibrium fares. Under this explanation, higher prices would reflect tighter capacity rather than changes in pricing behavior.

To assess this possibility, we directly examine whether capacity declines following AI adoption. Specifically, we regress the log number of departures and seats on AI adoption using the same baseline specification and fixed effects structure as in our main analysis and report the results in columns 4 and 5 of Table 11. Contrary to the capacity-based hypothesis, the estimates here indicate that both flight frequency and total seats tend to increase following enterprise AI adoption, which makes it unlikely that fare increases are driven by capacity reductions.

7.1.3 Revenue Management Optimization vs. Algorithmic Coordination

We next assess whether the observed price effects align more closely with revenue management optimization or algorithmic coordination. AI-aided revenue management systems could raise average fares by allocating seats more efficiently across fare buckets and smoothing pricing across transactions, which can result in lower price dispersion similar to that under price coordination. We perform two tests to provide suggestive evidence on these

competing explanations.

First, revenue management optimization can affect the level and distribution of fares, but it does not generate a clear prediction for how price dispersion should vary with market concentration. On the other hand, price coordination implies that incentives to undercut rivals are weaker in more concentrated routes, where fewer competitors and repeated interaction make deviations easier to detect and punish. As a result, price coordination becomes easier to sustain in high-HHI markets. In line with this intuition, Figure 3 provides descriptive evidence that more concentrated routes tend to have higher fares throughout the fare distribution. More formally, we test whether the compression of fare dispersion is stronger in more concentrated markets by interacting the number of airlines on a route that adopt enterprise AI pricing systems with route-level HHI and report the estimates in columns 1 to 3 of Table 5.[20]

These results, as discussed in Section 6.2, indicate that a higher route-level share of enterprise AI adoption is associated with lower fares on less concentrated routes, while the positive interaction with lagged HHI indicates that this negative effect becomes weaker and may turn positive as market concentration rises. The pattern holds for average fares as well as for both the lower and upper ends of the fare distribution. However, because the interaction effects are qualitatively similar across low-end and high-end fares, the estimates imply only modest changes in fare dispersion as concentration increases. Thus, although the results are consistent with coordination becoming easier to sustain in more concentrated markets, they provide only suggestive evidence of stronger fare compression in such markets.

Next, we examine whether low-end fares are realized in a more synchronized manner across carriers on a given route following AI adoption. Under revenue management optimization, price cuts reflect carrier-specific demand and inventory conditions, which implies that low-end fares should occur at different times across carriers and generate temporary dispersion. In contrast, coordination-based channels imply rapid matching of price cuts, so

[20]Columns 4–6 of Table 5 report qualitatively similar estimates based on GenAI platform use intensity.

that low-end fares occur at similar points in time and are quickly matched across airlines.

Figures 4 and 5 provide some corroborating descriptive evidence. Specifically, Figure 4 indicates that, for many routes, the share of low-fare observations is lower when all airlines on the route have adopted enterprise AI than when no airline has adopted it, as indicated by the concentration of points below the 45-degree line. At the same time, Figure 5 suggests that route-level fare dispersion does not fall systematically following enterprise AI adoption and may even increase for a substantial number of routes. When combined, these patterns do not appear to be supportive of a simple revenue management story in which carriers independently generate more low-fare realizations through idiosyncratic pricing adjustments. Instead, they are more consistent with the possibility that AI adoption is associated with fewer low-fare realizations without a corresponding decline in overall route-level fare variability.

In sum, our overall evidence thus far suggests a joint pattern of higher average fares, modestly more compressed price dispersion, stronger effects in more concentrated routes, and descriptive evidence consistent with fewer low-end fare realizations after AI adoption. This combination is difficult to reconcile with revenue management optimization alone and is instead more in line with tacit coordination facilitated by algorithmic pricing (Calvano et al., 2020).

7.2 Coordination Channels

The evidence presented in Section 7.1 appears most consistent with coordination-based channels. The GenAI ITT estimates in Table 4 corroborate this interpretation in that exposure to GenAI-enhanced pricing platforms on a route is associated with a sizable increase in average fares. Because $GenAI_{it}$ is an intention-to-treat indicator, the SVOI and VAII measures help further shed light on how these pricing outcomes map into specific coordination mechanisms, including the role of common algorithm vendors and the extent of platform diffusion within a route.

To provide some theoretical intuition, we develop a simple model in which GenAI platform use can affect equilibrium fares through common vendor overlap and platform diffusion. The model builds on the partial profit internalization framework commonly used in the industrial organization literature (e.g., Bresnahan 1982; Bernheim and Whinston 1990), where firms may place a positive weight on rivals' profits, which can capture coordinated conduct. We consider a differentiated Bertrand pricing game in which airlines on a route simultaneously choose prices, while vendor adoption influences both the set of rivals whose profits enter an airline's effective objective function and the overall intensity of platform use in the market. Under standard assumptions that guarantee single-valued best responses and strategic complementarity in prices, the pricing game is supermodular, which ensures the existence of a largest equilibrium price vector (Milgrom and Shannon, 1994; Topkis, 1998). To streamline the presentation, this section outlines the model setup and main takeaways, while more detailed discussion of the assumptions and derivations can be found in Online Appendix D.

Our model incorporates two distinct theoretical predictions that map directly to SVOI and VAII, our empirical measures of same-vendor overlap and platform diffusion, respectively. In particular, airline i's effective objective function is given by

$$\tilde{\pi}_i(\mathbf{p}; m) = \pi_i(\mathbf{p}; m) + \theta \sum_{i' \in S_i} \pi_{i'}(\mathbf{p}; m),$$

where S_i denotes the set of rival airlines on the route that share the same GenAI pricing vendor as airline i, $\mathbf{p}$ is a vector of fares in the same route, m is the intensity of vendor adoption in a route, and $\theta \in (0, 1)$ captures the degree of partial profit internalization among airlines using the same vendor. The model predicts that, as stated in Proposition 1, holding route-level adoption intensity m fixed, equilibrium prices weakly increase when vendor assignment expands same-vendor overlap. Specifically, for two vendor assignments v

and v' such that $S_i(v) \subseteq S_i(v')$,

$$\mathbf{p}^{\max}(m, v', \theta) \geq \mathbf{p}^{\max}(m, v, \theta).$$

This prediction aligns with the positive coefficient on SVOI in Table 4, which indicates that greater same-vendor overlap among competing airlines results in higher fares.

Second, the model distinguishes same-vendor overlap from overall platform diffusion within a route. This is captured by the route-level adoption intensity $m \in [0, 1]$, defined as the fraction of airlines on the route that use any GenAI pricing vendor. Assumption 3 imposes

$$\frac{\partial^2 \tilde{\pi}_i(\mathbf{p}; m)}{\partial p_i \partial m} \leq 0,$$

which states that the effective profit function exhibits decreasing differences in (p_i, m). Under Assumptions 1–3, the model further predicts that, as formalized in Proposition 2, the largest equilibrium price vector is weakly decreasing in m, so that for $m' \geq m$,

$$\mathbf{p}^{\max}(m', v, \theta) \leq \mathbf{p}^{\max}(m, v, \theta).$$

This prediction is in line with the negative coefficient on VAII in Table 4, which suggests that average fares tend to be lower on routes with wider GenAI platform diffusion.

The simple theoretical framework thus helps inform the directions of the estimates shown in Table 4. In the model, greater same-vendor overlap, captured empirically by higher SVOI, expands the set of rival airlines whose profits enter an airline's effective objective function, which dampens competition and raises equilibrium prices. This mechanism is also consistent with the partial profit internalization framework in Bresnahan (1982) and Bernheim and Whinston (1990) and the evidence in Aryal et al. (2024) that greater overlap in subcontracting relationships across carriers are associated with higher fares. In our setting, when rivals rely on the same GenAI-enhanced pricing vendor, pricing recommendations may

become more closely aligned, which can reduce the incentive to undercut and make parallel high-price outcomes easier to sustain.

On the other hand, our model also predicts that wider platform diffusion, or higher VAII, lowers equilibrium prices through the route-level adoption intensity term m, which reduces the return to further price increases. One possibility is that as GenAI-enhanced pricing platforms diffuse more broadly across carriers on a route, prices become more automated and strategically responsive, which can intensify competitive adjustment even when airlines rely on different vendors. This interpretation is consistent with the idea that algorithmic tools can accelerate reactions to rivals' posted fares and reduce pricing frictions, which can in turn strengthen competitive responses. More generally, more extensive diffusion does not necessarily imply stronger coordination; it can instead reflect a market environment in which wider platform diffusion leads airlines to respond more quickly to rivals' prices, which can lower fares.

As discussed in Section 2.3, the ex ante effects of same-vendor overlap and platform diffusion are both theoretically ambiguous. Shared vendor adoption could either raise fares by aligning pricing incentives across rivals or lower fares if vendor self-competition and rapid algorithmic reactions trigger stronger undercutting (Harrington, 2022; Musolff, 2025). Likewise, broader diffusion of pricing platforms could either facilitate coordination by making rivals' behavior more predictable (Calvano et al., 2020; Wang et al., 2024b) or lower fares by improving the speed and precision of competitive responses across the market. Our estimates indicate that, in the data, the first force tends to dominate for same-vendor overlap, while the second appears to dominate for platform diffusion. Our theoretical framework thus offers a more structured interpretation of why SVOI and VAII are associated with opposite average-fare effects in Table 4.

7.3 Heterogeneity by Route Size and Baseline Scale

Airlines do not disclose the specific routes or flights on which GenAI-enhanced pricing tools are deployed, and public statements generally emphasize that these systems are integrated into broader revenue management platforms rather than activated on a route-by-route basis. As a result, one cannot directly observe route-level implementation of GenAI pricing tools. Disclosures by airlines and vendors also confirm gradual diffusion of pricing platforms that incorporate GenAI components, rather than discrete and publicly identifiable "GenAI routes." For example, Delta has stated that its GenAI-enhanced pricing system was initially implemented on a small share of its domestic network and would expand over time. Vendor materials likewise emphasize gradual scaling and experimentation, including A/B testing, rather than fixed deployment on a clearly identifiable set of routes.[21] This diffusion-based rollout implies that airline-level adoption need not translate into a uniform or directly observable route-level treatment.

Accordingly, we do not interpret route characteristics as proxies for whether GenAI is deployed on a given route. Instead, we examine whether the pricing effects of airline-level AI adoption are stronger on routes where these technologies are likely to have larger observable effects on pricing outcomes. In particular, higher-traffic routes involve greater passenger volume and a larger scale over which pricing systems can affect market outcomes. This motivates two complementary heterogeneity analyses based on predetermined route size.

First, Table 12 divides routes into small and large markets based on the average number of tickets sold in 2013–2014, that is, before any airline in our sample adopted enterprise AI. The results indicate that the pricing effects of both enterprise AI adoption and GenAI platform use are stronger on larger routes. For enterprise AI, the estimated effect on average fares is close to zero on small routes, but significantly positive on large routes. A similar contrast appears across the fare distribution: the lower-end and upper-end fare effects are both more

[21] See Fetcherr, "Dynamic Pricing in Aviation: How AI Is Revolutionizing Airline Revenue Management," company blog, December 5, 2024, which describes implementation through "Gradual Scaling and A/B Testing" as part of its rollout process, `https://www.fetcherr.io/blog/dynamic-pricing-in-aviation`.

positive on large routes, whereas the small-route estimates are weaker and somewhat mixed. For GenAI platform use, the estimated effects are positive throughout the fare distribution in both subsamples, but they are considerably larger on large routes. Specifically, average fares rise by about 3.6% on small routes and 6.3% on large routes, with similar differences at the 20th and 80th percentiles. Overall, Table 12 suggests that AI-related pricing effects are more pronounced in high-traffic markets, where algorithmic pricing is likely to operate at greater scale.

Next, we estimate a continuous baseline scale specification for GenAI platform use that interacts the GenAI use indicator with predetermined route passenger volume, $Passengers_j^{pre}$. Specifically,

$$BaselineScale_j^{pre} = \log\left(1 + Passengers_j^{pre}\right),$$

and $Passengers_j^{pre}$ denotes total route-level passengers measured before 2023Q4, when the first GenAI-enhanced pricing platform was adopted in our sample. Because $BaselineScale_j^{pre}$ is constructed entirely from pre-adoption data, it reflects predetermined route traffic volume rather than post-adoption market outcomes. This also makes the measure less vulnerable to reverse causality than contemporaneous fares, OTP, or realized revenue. In addition, the baseline scale measure is directly aligned with our diffusion-based interpretation of GenAI adoption in that platform-based pricing should be most salient on routes with substantial traffic volume, where pricing systems operate at greater scale.

Table 13 presents the corresponding estimates. The interaction term between *GenAI* and the predetermined baseline scale is positive and precisely estimated for average fares as well as for lower-end and upper-end fares. This implies that the effect of GenAI platform use becomes more positive as route baseline scale increases. Although the baseline coefficient on *GenAI* is negative, it reflects the effect at very low baseline scale levels; when evaluated at the sample mean of the baseline scale, the overall effect remains positive. The estimates therefore reinforce the evidence from the subsample analysis in Table 12: conditional on airline-level GenAI platform use, the observed fare effects are stronger on routes with greater

pre-adoption traffic volume.

8 Welfare Implications

Our results indicate that enterprise AI adoption generates dual effects by enhancing operational efficiency through improved on-time performance and air quality, while also raising average fares and compressing price dispersion. To quantify the net welfare impact on consumers, we estimate a simple logit demand model for airline tickets that includes airfares, delay time, and AQI. In particular, we first specify the indirect utility of a passenger k choosing airline i on route j at time t as:

$$u_{ijkt} = \alpha p_{ijt} + \beta_1 \text{Delay}_{ijt} + \beta_2 \text{AQI}_{jt} + \mathbb{X}'_{ijt}\theta + \mu_i + \lambda_j + \delta_t + \epsilon_{ijkt},$$

where $\mathbb{X}_{ijt}$ denotes airline and route characteristics and ϵ_{ijkt} captures unobserved consumer heterogeneity. Assuming that ϵ_{ijkt} follows a Type I extreme value distribution, we follow Berry (1994) and estimate the corresponding demand function as:

$$\ln\left(\frac{s_{ijt}}{s_{0jt}}\right) = \alpha p_{ijt} + \beta_1 \text{Delay}_{ijt} + \beta_2 \text{AQI}_{jt} + \mathbb{X}'_{ijt}\theta + \mu_i + \lambda_j + \delta_t,$$

where the market size is defined as the sum of origin and destination county populations and s_{0jt} denotes the outside-option market share. Since delay time and AQI are likely to be higher at airports with greater passenger volume, and unobserved factors may simultaneously affect delay time, AQI, and market share, we use instrumental variables for airfares, delay time, and AQI. The IVs employed in the demand regressions include the sum of competitors' characteristics on the same route and airline i's average route distance in other markets.

Following the market share inversion method of Berry (1994), we compute the welfare effects by comparing the predicted mean utility under observed outcomes with the counterfactual mean utility in the absence of AI adoption. We use estimates from Tables 2 and 3 and

assume that, without AI adoption, the average fare would be 1.01% lower, delay time 0.39 minutes longer, and AQI 2.16 units higher. For the price effect, we compute predicted mean utility $\hat{v}_{ijt}$ and mean utility evaluated at the counterfactual fare without AI adoption, $\hat{v}_{ijt}^{cf}$, i.e., $\hat{v}_{ijt}\,(p_{ijt}/1.0101)$. We then divide the difference between the two utilities by the marginal utility of income (as captured by the estimated coefficient $\hat{\alpha}$) to express the utility difference in monetary terms. The calculation for the consumer surplus gain from improved on-time performance follows a similar approach. In particular, we compute the counterfactual utility by increasing delay time by 0.39 minutes for airline i and route j in periods after AI adoption. Likewise, since Table 2 shows that AQI decreases by 2.16 following AI adoption, we add 2.16 to post-adoption observations to compute the corresponding counterfactual utility.

Overall, our calculations indicate substantial welfare gains from improved service quality of approximately $42.9 million and price-induced welfare losses of about $3.1 million, implying a net gain in consumer surplus. The estimated environmental benefits amount to $0.7 million. These estimates thus suggest that AI adoption in the airline industry increases overall welfare through improved product quality and reduced negative environmental externalities, despite modest price increases.

9 Conclusion

Leveraging empirical evidence from the U.S. airline industry, our study finds that enterprise AI adoption improves product quality and local air quality while raising average fares and compressing price dispersion. For generative AI–enhanced pricing tools, the price effects become larger and extend across the fare distribution, but their direction and magnitude depend on the extent and pattern of diffusion across carriers. Specifically, when rivals rely on the same algorithm vendor, fares tend to be higher, especially at the lower end of the distribution, whereas the diffusion of GenAI-enhanced pricing platforms is associated with lower fares across the distribution, with the effect again being more pronounced at the lower

end.

Our findings raise important questions for competition policy. As AI-driven pricing becomes more widespread, regulatory frameworks need to address the potential for unintended coordination among algorithmic platforms. Policymakers have recently flagged concerns about the use of AI to personalize airfares for individual consumers, which shows that, beyond route-level coordination, discriminatory pricing at the individual customer level is also under scrutiny.[22] Our results suggest that the risks of algorithmic coordination depend less on the mere presence of AI-driven pricing tools than on whether rivals rely on a common vendor and how these platforms diffuse across competing firms. That is, not all AI adoption is equal, so regulatory oversight should focus not only on whether AI-driven tools are used but also how they are deployed across firms and markets.

Our study also opens several avenues for future research. First, because AI adoption in the airline industry is still relatively recent, our analysis focuses on short-term effects; in particular, the GenAI estimates primarily capture outcomes among early users rather than a fully diffused pricing regime. While this provides insight into the immediate impact of AI, assessing longer-term effects will require additional data as platform diffusion progresses. Second, several carriers, most notably American Airlines, have begun developing their own generative AI systems. As these in-house pricing tools are rolled out, the composition of pricing technologies across routes will change, which can in turn affect how platform diffusion is measured and interpreted. Future work should therefore revisit and adapt the VAII measure to account for these developments. Moreover, access to more granular data on the timing and scope of GenAI platform implementation across routes would allow for a more nuanced assessment of how gradual rollout and diffusion affect market outcomes and represent another fruitful direction for future research.

[22]Source: David Shepardson, "U.S. criticizes use of AI to personalize airline ticket prices, would investigate," *Reuters*, August 5, 2025, `https://www.reuters.com/world/us/us-criticizes-use-ai-personalize-airline-ticket-prices-would-investigate-2025-08-05/`.

References

Aryal, G., Campbell, D. J., Ciliberto, F., and Khmelnitskaya, E. A. (2024). "Common subcontracting and airline prices." *Review of Economics and Statistics*, Advance online publication, https://doi.org/10.1162/rest_a_01420.

Assad, S., Clark, R., Ershov, D., and Xu, L. (2024). "Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market." *Journal of Political Economy*, *132*(3), 723–771.

Azar, J., and Ribeiro, R. (2021). "Estimating oligopoly with shareholder voting models." *SSRN Electronic Journal*, IESE Business School Working Paper.

Azar, J., Schmalz, M. C., and Tecu, I. (2018). "Anticompetitive effects of common ownership." *Journal of Finance*, *73*(4), 1513–1565.

Azoulay, P., Krieger, J., and Nagaraj, A. (2025). "Old moats for new models: Openness, control, and competition in generative artificial intelligence." *Entrepreneurship and Innovation Policy and the Economy*, *4*(1), 7–46.

Bai, J. (2009). "Panel data models with interactive fixed effects." *Econometrica*, *77*(4), 1229–1279.

Bernheim, B. D., and Whinston, M. D. (1990). "Multimarket contact and collusive behavior." *The RAND Journal of Economics*, *21*(1), 1–26.

Berry, S. T. (1994). "Estimating discrete-choice models of product differentiation." *RAND Journal of Economics*, *25*(2), 242–262.

Borenstein, S., and Rose, N. L. (1994). "Competition and price dispersion in the US airline industry." *Journal of Political Economy*, *102*(4), 653–683.

Bresnahan, T. F. (1982). "The oligopoly solution concept is identified." *Economics Letters*, *10*(1), 87–92.

Callaway, B., and Sant'Anna, P. H. (2021). "Difference-in-differences with multiple time periods." *Journal of Econometrics*, *225*(2), 200–230.

Calvano, E., Calzolari, G., Denicolo, V., and Pastorello, S. (2020). "Artificial intelligence, algorithmic pricing, and collusion." *American Economic Review*, *110*(10), 3267–3297.

Chávez-Miranda, E., Martínez-Torres, M. d. R., Palma, L., and López-Escobar, A. (2025). "Can AI-powered pricing systems collude without human intervention? Evidence from the hotel industry." Working paper.

Ge, Q., Kim, D., Kim, M., and Yoon, J. (2025a). "Network structures and resilience of supply: Insights from the US airline industry." Working paper.

Ge, Q., Kim, M., and Rupp, N. (2025b). "Algorithmic collusion in the skies: The role of AI in shaping airline competition." *CPI Antitrust Chronicle*, April, Vol. 1, 22–28.

Gerardi, K. S., and Shapiro, A. H. (2009). "Does competition reduce price dispersion? New evidence from the airline industry." *Journal of Political Economy*, *117*(1), 1–37.

Geske, A. M., Herold, D. M., and Kummer, S. (2024). "Artificial intelligence as a driver of efficiency in air passenger transport: A systematic literature review and future research avenues." *Journal of the Air Transport Research Society*, *3*(0), 100030.

Goldberg, S., and Lam, H. T. (2025). "Generative AI in equilibrium: Evidence from a creative goods marketplace." Working paper.

Goodman-Bacon, A. (2021). "Difference-in-differences with variation in treatment timing." *Journal of Econometrics*, *225*(2), 254–277.

Gu, Y. (2023). "Can a dynamic pricing algorithm facilitate tacit collusion? Evidence from the airline industry." Working paper.

Han, K., and Mithas, S. (2013). "Information technology outsourcing and non-IT operating costs: An empirical investigation." *MIS Quarterly*, 315–331.

Harrington, J. E. (2022). "The effect of outsourcing pricing algorithms on market competition." *Management Science*, *68*(9), 6889–6906.

Harrington, J. E. (2025a). "Cost coordination." *RAND Journal of Economics*, Advance online publication, `https://doi.org/10.1111/1756-2171.70000`.

Harrington, J. E. (2025b). "An economic test for an unlawful agreement to adopt a third-party's pricing algorithm." *Economic Policy*, *40*(121), 261–295.

Hortaçsu, A., Natan, O. R., Parsley, H., Schwieg, T., and Williams, K. R. (2024). "Organizational structure and pricing: Evidence from a large US airline." *Quarterly Journal of Economics*, *139*(2), 1149–1199.

Kim, M., Shen, L., and Ge, Q. (2023). "Does competition increase or decrease price dispersion? Insights from one-way vs. round-trip airfares." *Oxford Bulletin of Economics and Statistics*, *85*(2), 435–455.

Kühn, K.-U., and Tadelis, S. (2018). "The economics of algorithmic pricing: Is collusion really inevitable?" Working paper.

Mazzeo, M. J. (2003). "Competition and service quality in the US airline industry." *Review of Industrial Organization*, *22*(4), 275–296.

Milgrom, P., and Shannon, C. (1994). "Monotone comparative statics." *Econometrica*, *62*(1), 157–180.

Musolff, L. (2025). "Algorithmic pricing, price wars and tacit collusion: Evidence from e-commerce." Working paper.

Prince, J. T., and Simon, D. H. (2015). "Do incumbents improve service quality in response to entry? Evidence from airlines' on-time performance." *Management Science*, *61*(2), 372–390.

Sun, L., and Abraham, S. (2021). "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of Econometrics*, *225*(2), 175–199.

Sweeting, A., Roberts, J. W., and Gedge, C. (2020). "A model of dynamic limit pricing with an application to the airline industry." *Journal of Political Economy*, *128*(3), 1148–1193.

Tafur, C. L., Camero, R. G., Rodríguez, D. A., Rincón, J. C. D., and Saenz, E. R. (2025). "Applications of artificial intelligence in air operations: A systematic review." *Results in Engineering*, *25*, 103742.

Topkis, D. M. (1998). *Supermodularity and Complementarity*. Princeton University Press.

US Environmental Protection Agency (2025). "Air data: Air quality data collected at outdoor monitors across the US." Last updated September 2, 2025.

Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., et al. (2024a). "A survey on large language model based autonomous agents." *Frontiers of Computer Science*, *18*(6), 186345.

Wang, Q., Huang, Y., Singh, P. V., and Srinivasan, K. (2024b). "Algorithms, artificial intelligence and simple rule-based pricing." Working paper.

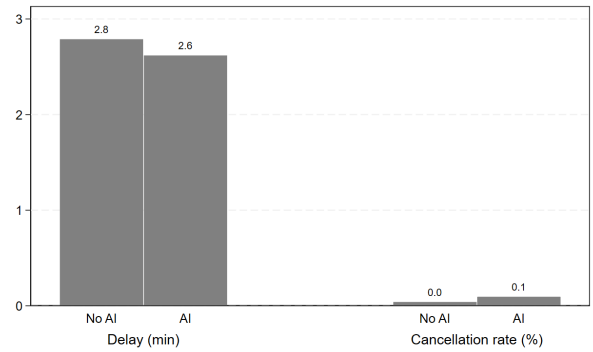
Weyl, E. G., and Fabinger, M. (2013). "Pass-through as an economic tool: Principles of incidence under imperfect competition." *Journal of Political Economy*, *121*(3), 528–583.

Zhang, S., Mehta, N., Singh, P. V., and Srinivasan, K. (2021). "Frontiers: Can an artificial intelligence algorithm mitigate racial economic inequality? An analysis in the context of Airbnb." *Marketing Science*, *40*(5), 813–820.

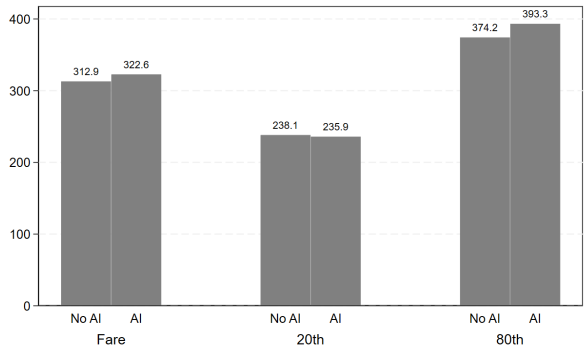
Zhao, H., and Berman, R. (2025). "Algorithmic collusion of pricing and advertising on e-commerce platforms." Working paper.

Figure 1: AI Adoption, Service Quality, Fares, and Vendor Indices

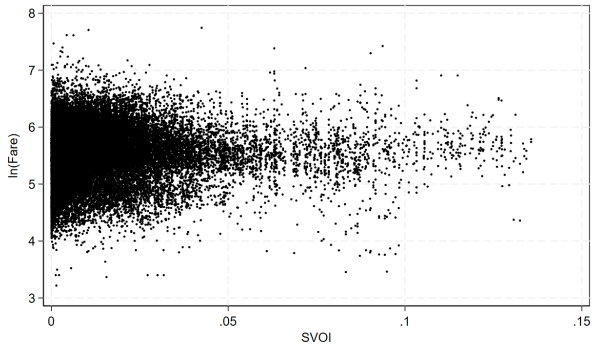
(a) OTP by Enterprise AI Adoption



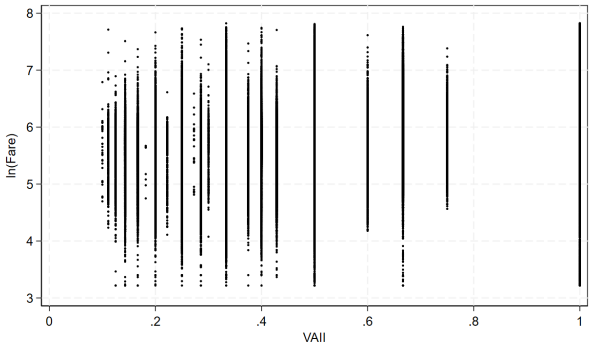
(b) Fare by Enterprise AI Adoption



(c) SVOI vs. Fares



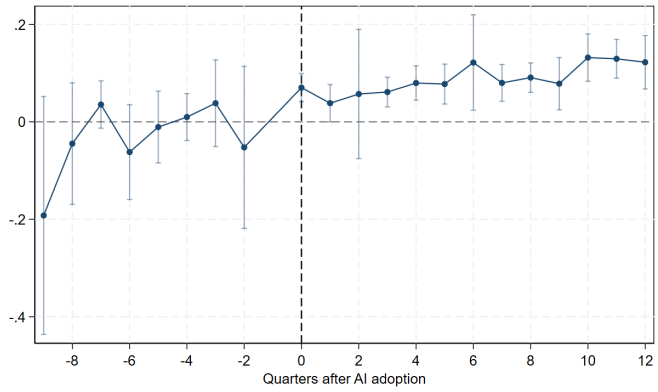
(d) VAII vs. Fares



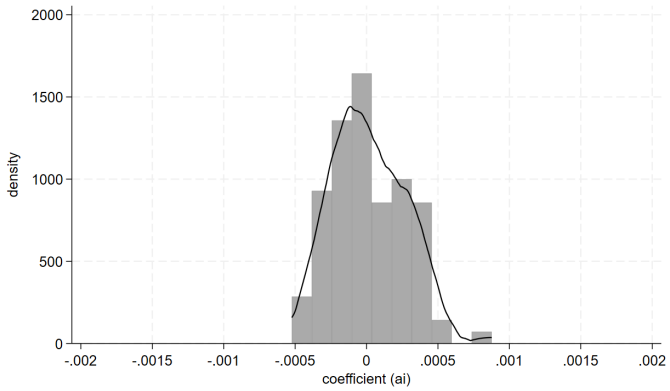
Note: Panels (c) and (d) include only observations with positive SVOI and VAII values.

Figure 2: Robustness check: Staggered Adoption and Placebo Test

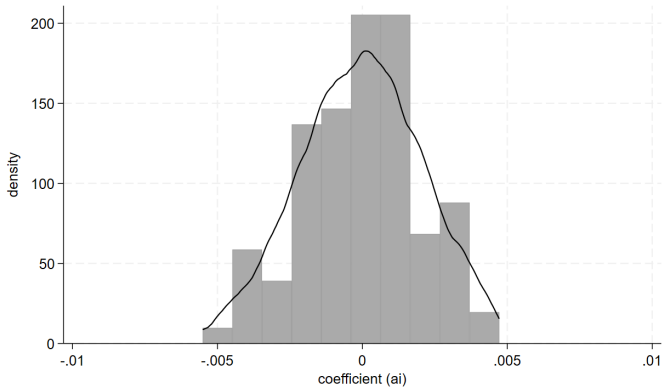
(a) Dynamic Impact of Enterprise AI Adoption on Airfare



(b) Placebo Test for Enterprise AI



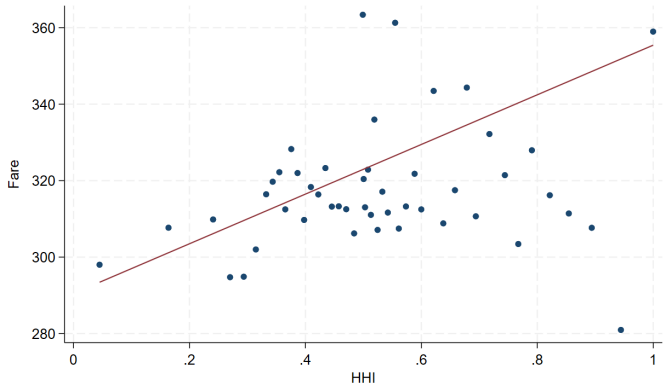
(c) Placebo Test for GenAI



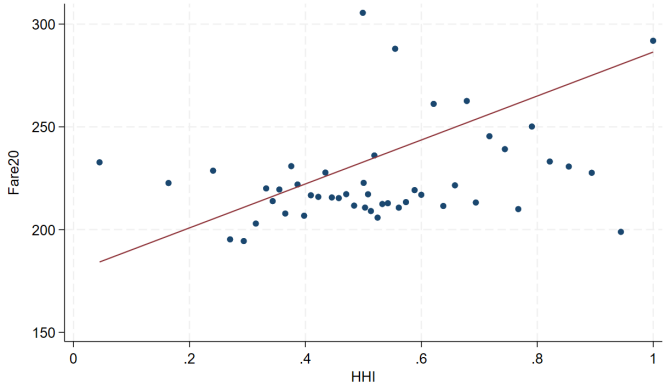
Note: Panel (a) restricts the sample to periods through 2021 and thus excludes the post-2021 period when all airlines had adopted enterprise AI. In the first stage, lagged HHI and MMC are regressed on the lagged number of LCCs, the sum of competitors' prices, passengers, and distances in the market. The corresponding fitted values are then used in the second-stage regression. Panels (b) and (c) report the distributions of estimated coefficients from separate placebo tests based on randomly reassigned enterprise AI adoption and GenAI platform use, respectively.

Figure 3: Relationship Between HHI and Fares

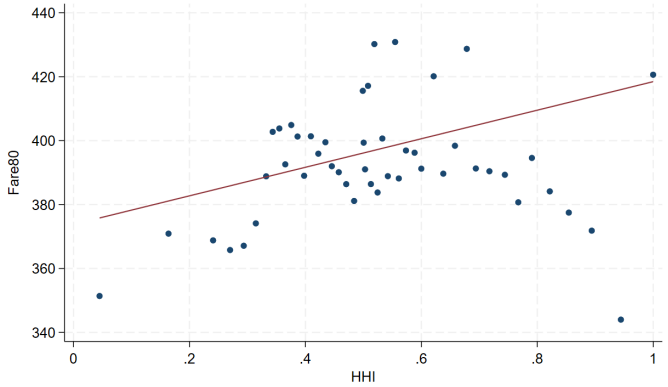
(a) lnFare



(b) lnFare20

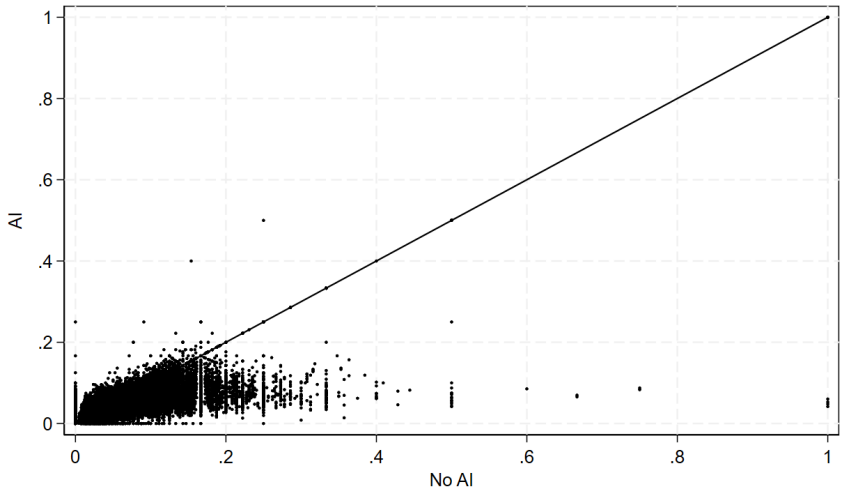


(c) lnFare80



Note: Each panel shows a binned scatter plot of a fare measure against HHI, reported separately for average fares, 20th percentile fares, and 80th percentile fares. The HHI distribution is divided into equal-sized bins, and the mean fare within each bin is plotted. A fitted linear regression line is overlaid to illustrate the overall relationship between HHI and each fare measure.

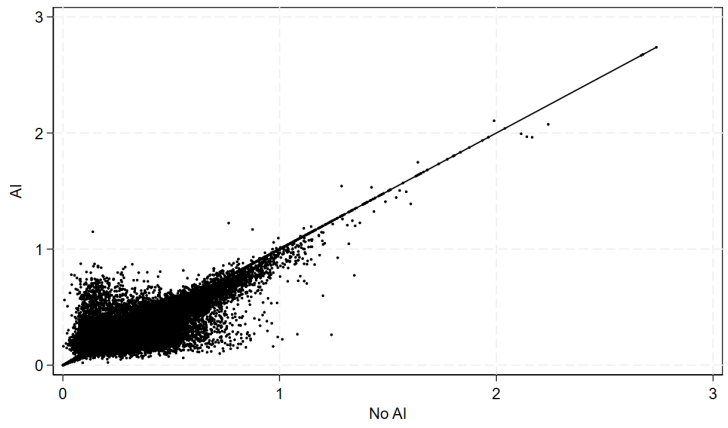
Figure 4: Share of Low Fares by Enterprise AI Adoption



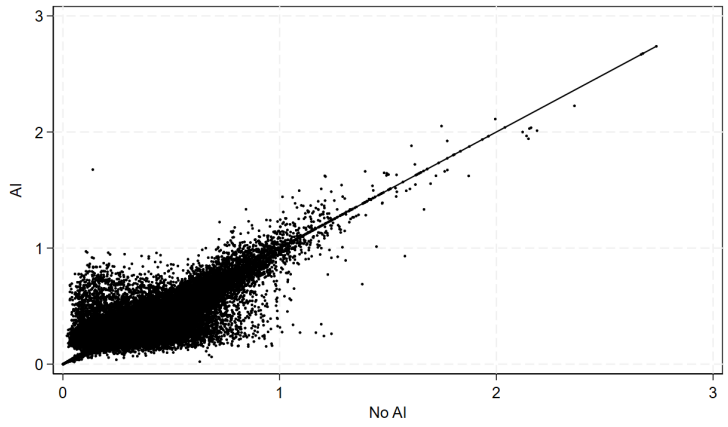
Note: The figure compares the share of low fares on a route during periods when no airline had adopted enterprise AI with the share during periods when all airlines on the route had adopted. For each route, the 10th percentile of fares is computed using the full sample period, and a low fare is defined as any fare below that threshold. The average share of low-fare observations is then calculated separately for the no-AI period and the all-AI period. The solid line is the 45-degree line; points below it indicate that the share of low fares is lower in the all-AI period than in the no-AI period.

Figure 5: Standard Deviation of Fares by Enterprise AI Adoption

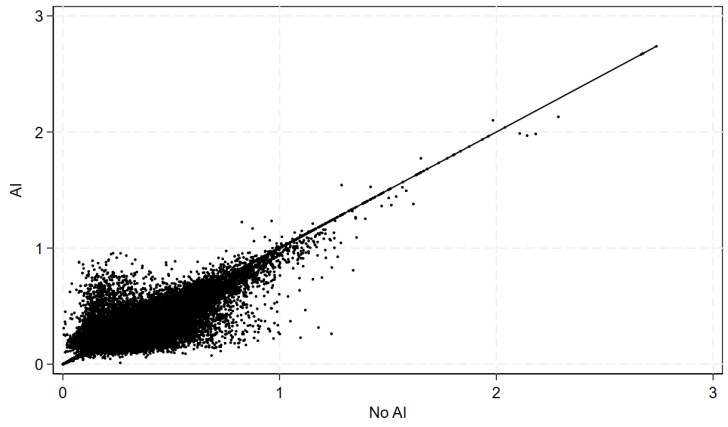
(a) lnFare



(b) lnFare20



(c) lnFare80



Note: Each panel compares the standard deviation of a route-level passenger-weighted fare measure between periods when no airline on the route had adopted enterprise AI and periods when all airlines on the route had adopted. The three panels report these comparisons separately for average fares, 20th percentile fares, and 80th percentile fares. The solid line is the 45-degree line; points above it indicate that fare variability is higher in the all-AI period than in the no-AI period.

Table 1: Summary Statistics

	Enterprise AI Adoption			GenAI Platform Use		
	Full sample	Not Adopted	Adopted	Full sample	Not Adopted	Adopted
	(1)	(2)	(3)	(4)	(5)	(6)
Enterprise AI Adoption	0.715	0	1	0.763	0.744	1
	(0.452)	(0)	(0)	(0.425)	(0.437)	(0)
GenAI Platform Use				0.0745	0	1
				(0.263)	(0)	(0)
SVOI				0.0000907	0.0000360	0.000771
				(0.00178)	(0.000977)	(0.00551)
VAII				0.0706	0.0235	0.655
				(0.206)	(0.104)	(0.250)
Fare	310.4	301.7	313.8	325.6	321.1	381.5
	(138.5)	(136.1)	(139.2)	(149.3)	(146.5)	(171.8)
20th Pct Fare	225.4	225.3	225.5	238.8	236.4	268.2
	(129.9)	(128.9)	(130.3)	(141.6)	(139.2)	(166.2)
80th Pct Fare	378.7	363.3	384.8	396.3	390.2	471.1
	(178.4)	(172.9)	(180.2)	(190.1)	(186.0)	(221.8)
Air Quality Index	22.55	22.39	22.61			
	(17.70)	(18.62)	(17.32)			
Delay Time	3.052	3.303	2.952			
	(15.44)	(14.84)	(15.67)			
Cancellation Rate (%)	0.0901	0.0462	0.108			
	(1.661)	(0.980)	(1.864)			
Sample Period	2013Q1–2023Q3			2013Q1–2025Q2		
Observations	3,697,957	1,055,346	2,642,611	5,488,575	5,079,504	409,071

Table 2: Impact of Enterprise AI Adoption on On-Time Performance and Air Quality

	(1) Delay Time	(2) Cancellation (%)	(3) AQI
AI Adoption	-0.394***	0.070***	-2.161***
	(0.056)	(0.004)	(0.048)
Observations	3,304,316	3,304,316	2,682,106
	(4) Delay Time	(5) Cancellation (%)	(6) AQI
Pr(AI $\widehat{\text{Adoption}}$)	-0.398***	0.079***	-2.220***
	(0.057)	(0.004)	(0.049)
Observations	3,304,316	3,304,316	2,682,106
Controls	Y	Y	Y
Airline FE	Y	Y	Y
Route FE	Y	Y	Y
Time FE	Y	Y	Y

Note: Sample period is 2013Q1–2023Q3. *AQI* is the average AQI weighted by fuel expenses multiplied by passenger shares. *Pr(AI $\widehat{\text{adoption}}$)* is the predicted probability of AI adoption from the first stage of the generalized regressor approach. The first-stage results are reported in Appendix Table A1. Control variables include lagged HHI, MMC, distance and its quadratic term, lagged origin and destination network size, the lagged number of tickets sold, and a merger dummy which is 1 if airline experienced merger at time t, and 0 otherwise. Instrumental variables for the lag HHI and MMC are the lagged number of LCCs in each route and the sum of competitors' prices, passengers, and distances in the market. Standard errors in parentheses are clustered by route. *** p<0.01, ** p<0.05, * p<0.1.

Table 3: Impact of Enterprise AI Adoption on Airfare Distribution

	(1) lnFare	(2) lnFare20	(3) lnFare80
AI Adoption	0.010***	0.025***	0.008***
	(0.001)	(0.001)	(0.001)
Observations	3,304,316	3,304,316	3,304,316
	(4) lnFare	(5) lnFare20	(6) lnFare80
$\widehat{\text{Pr(AI Adoption)}}$	0.010***	0.024***	0.007***
	(0.001)	(0.001)	(0.001)
Observations	3,304,316	3,304,316	3,304,316
Controls	Y	Y	Y
Airline FE	Y	Y	Y
Route FE	Y	Y	Y
Time FE	Y	Y	Y

Note: Sample period is 2013Q1–2023Q3. *$\widehat{Pr(AI\ adoption)}$* is the predicted probability of AI adoption from the first stage of the generalized regressor approach. The first-stage results reported in Appendix Table A1. Control variables include lagged HHI, MMC, distance and its quadratic term, lagged origin and destination network size, the lagged number of tickets sold, and a merger dummy which is 1 if airline experienced merger at time t, and 0 otherwise. Instrumental variables for the lag HHI and MMC are the lagged number of LCCs in each route and the sum of competitors' prices, passengers, and distances in the market. Standard errors in parentheses are clustered by route. *** p<0.01, ** p<0.05, * p<0.1.

Table 4: Impact of GenAI Platform Use and Vendor Sharing on Airfare Distribution

	(1) lnFare	(2) lnFare20	(3) lnFare80
GenAI	0.050***	0.037***	0.041***
	(0.001)	(0.001)	(0.001)
Observations	3,687,969	3,687,969	3,687,969
	(4) lnFare	(5) lnFare20	(6) lnFare80
GenAI	0.077***	0.039***	0.063***
	(0.002)	(0.003)	(0.003)
SVOI	0.430**	1.945***	0.744***
	(0.176)	(0.221)	(0.202)
VAII	-0.156***	-0.275***	-0.134***
	(0.004)	(0.005)	(0.005)
GenAI × SVOI	0.209	-1.714***	-0.660***
	(0.185)	(0.222)	(0.193)
GenAI × VAII	-0.025***	0.044***	-0.020***
	(0.005)	(0.006)	(0.005)
Marginal Effects			
GenAI	0.075***	0.042***	0.062***
	(0.002)	(0.003)	(0.002)
SVOI	0.446***	1.812***	0.693***
	(0.164)	(0.208)	(0.190)
VAII	-0.158***	-0.272***	-0.136***
	(0.004)	(0.005)	(0.005)
Observations	3,687,969	3,687,969	3,687,969
Controls	Y	Y	Y
Airline FE	Y	Y	Y
Route FE	Y	Y	Y
Time FE	Y	Y	Y

Note: Sample period is 2015Q1–2025Q2. The sample includes Delta, United, American, Frontier, Southwest, and JetBlue. See Table 3 for other notes. Standard errors in parentheses are clustered by route. *** p<0.01, ** p<0.05, * p<0.1.

Table 5: Robustness Check: Impact of Enterprise AI Adoption and GenAI Platform Use (Route-Level Estimates)

Enterprise AI	(1) lnFare	(2) lnFare20	(3) lnFare80
AI Share	-1.490***	-1.878***	-1.810***
	(0.154)	(0.162)	(0.213)
AI Share $\widehat{\times}$ Lag HHI	2.178***	2.808***	2.630***
	(0.259)	(0.273)	(0.341)
Observations	1,268,530	1,268,530	1,478,449
Controls	Y	Y	Y
Route FE	Y	Y	Y
Time FE	Y	Y	Y
GenAI	(4) lnFare	(5) lnFare20	(6) lnFare80
GenAI Share	-2.343***	-3.733***	-3.874***
	(0.595)	(0.732)	(0.834)
GenAI Share $\widehat{\times}$ Lag HHI	3.485***	5.586***	5.870***
	(0.949)	(1.173)	(1.334)
Observations	1,470,572	1,470,572	1,470,572
Controls	Y	Y	Y
Route FE	Y	Y	Y
Time FE	Y	Y	Y

Note: Sample period is 2013Q1–2023Q3 for columns 1–3 and 2013Q1–2025Q2 for columns 4–6. *AI Share* is the fraction of airlines that adopted the AI in each route-time. *GenAI Share* is the fraction of airlines that adopted the GenAI pricing in each route-time. Control variables are average of each variable in each route. Instrumental variables for lagged HHI, MMC, AI Share×HHI, and GenAI Share×HHI are administrative expenses, reservation expenses, and line service expenses multiplied by total number of tickets sold in each route. Standard errors in parentheses are clustered by airline. *** p<0.01, ** p<0.05, * p<0.1.

Table 6: Robustness Check: Impact of GenAI Platform Use and Vendor Sharing on Airfare Distribution (Control Function Approach)

	(1) lnFare	(2) lnFare20	(3) lnFare80
GenAI	0.066***	0.060***	0.054***
	(0.001)	(0.001)	(0.001)
SVOI	7.038***	7.479***	4.752***
	(0.448)	(0.518)	(0.378)
VAII	-0.435***	-0.454***	-0.315***
	(0.012)	(0.015)	(0.012)
Observations	3,687,969	3,687,969	3,687,969
Controls	Y	Y	Y
Airline FE	Y	Y	Y
Route FE	Y	Y	Y
Time FE	Y	Y	Y

Note: The sample includes Delta, United, American, Frontier, Southwest, and JetBlue. Residuals from the first-stage regressions for SVOI and VAII are included as additional controls in the second stage; their coefficients are statistically significant, which is consistent with potential endogeneity of SVOI and VAII. See Table 3 for other notes. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 7: Robustness Check: Impact of Enterprise AI Adoption on Airfare Distribution (Alternative Fixed Effects)

	(1) lnFare	(2) lnFare20	(3) lnFare80
AI Adoption	0.011***	0.026***	0.009***
	(0.001)	(0.001)	(0.001)
Observations	3,304,316	3,304,316	3,304,316
Controls	Y	Y	Y
Airline FE	Y	Y	Y
Route FE	Y	Y	Y
Route FE × Trend	Y	Y	Y
Time FE	Y	Y	Y

	(4) lnFare	(5) lnFare20	(6) lnFare80
AI Adoption	0.023***	0.036***	0.019***
	(0.001)	(0.001)	(0.001)
Observations	3,304,316	3,304,316	3,304,316
Controls	Y	Y	Y
Airline FE	Y	Y	Y
Airline FE × Trend	Y	Y	Y
Route FE	Y	Y	Y
Time FE	Y	Y	Y

	(7) lnFare	(8) lnFare20	(9) lnFare80
AI Adoption	0.018***	0.022***	0.018***
	(0.001)	(0.001)	(0.001)
Observations	3,296,395	3,296,395	3,296,395
Controls	Y	Y	Y
Airline × Route FE	Y	Y	Y
Route FE	Y	Y	Y
Time FE	Y	Y	Y

Note: Sample period is 2013Q1–2023Q3. See Table 3 for other notes. *** p<0.01, ** p<0.05, * p<0.1.

Table 8: Robustness Check: Impact of GenAI Platform Use and Vendor Sharing on Airfare Distribution (Alternative Fixed Effects)

	(1) lnFare	(2) lnFare20	(3) lnFare80
GenAI	0.059***	0.035***	0.042***
	(0.002)	(0.003)	(0.003)
SVOI	0.777***	1.429***	0.752***
	(0.176)	(0.221)	(0.204)
VAII	-0.151***	-0.274***	-0.129***
	(0.004)	(0.005)	(0.005)
GenAI × SVOI	-0.163	-0.979***	-0.558***
	(0.186)	(0.231)	(0.206)
GenAI × VAII	-0.035***	0.024***	-0.025***
	(0.005)	(0.006)	(0.005)
Marginal Effects			
GenAI	0.057***	0.037***	0.040***
	(0.002)	(0.003)	(0.002)
SVOI	0.764***	1.353***	0.708***
	(0.164)	(0.207)	(0.192)
VAII	-0.154***	-0.272***	-0.131***
	(0.004)	(0.005)	(0.005)
Observations	3,681,269	3,681,269	3,681,269
Controls	Y	Y	Y
Airline × Route FE	Y	Y	Y
Route FE	Y	Y	Y
Time FE	Y	Y	Y

Note: Sample period is 2013Q1–2025Q2. The sample includes Delta, United, American, Frontier, Southwest, and JetBlue. See Table 3 for other notes. *** p<0.01, ** p<0.05, * p<0.1.

Table 9: Robustness Check: Impact of Enterprise AI Adoption and GenAI Platform Use on Airfare Distribution (Interactive Factor Model)

Enterprise AI	(1) lnFare	(2) lnFare20	(3) lnFare80
AI adoption	0.018***	0.025***	0.016***
	(0.001)	(0.001)	(0.001)
Observations	623,448	623,448	623,448
Controls	Y	Y	Y
Airline FE	Y	Y	Y
Route FE	Y	Y	Y
Time FE	Y	Y	Y
GenAI	(4) lnFare	(5) lnFare20	(6) lnFare80
GenAI	0.021***	0.007***	0.001
	(0.001)	(0.002)	(0.001)
Observations	671,111	671,111	671,111
Controls	Y	Y	Y
Airline FE	Y	Y	Y
Route FE	Y	Y	Y
Time FE	Y	Y	Y

Note: Sample period is 2013Q1–2023Q3 for columns 1–3 and 2013Q1–2025Q2 for columns 4–6. To reduce computational burden, the sample is restricted to routes with more than 1,000 passengers for columns 1–3 and more than 500 passengers for columns 4–6. All specifications are estimated using the interactive fixed-effects model of Bai (2009). Control variables include lagged HHI, MMC, distance and its quadratic term, lagged origin and destination network size, and the lagged number of tickets sold. *** p<0.01, ** p<0.05, * p<0.1.

Table 10: Robustness Check: Leave-One-Airline-Out Regression for GenAI Platform Use

	AA			B6			DL		
	(1) lnFare	(2) lnFare20	(3) lnFare80	(4) lnFare	(5) lnFare20	(6) lnFare80	(7) lnFare	(8) lnFare20	(9) lnFare80
GenAI	0.043***	-0.019***	0.038***	0.052***	0.039***	0.042***	0.051***	0.044***	0.047***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	2,576,542	2,576,542	2,576,542	3,619,455	3,619,455	3,619,455	2,583,816	2,583,816	2,583,816
	F9			UA			WN		
	(10) lnFare	(11) lnFare20	(12) lnFare80	(13) lnFare	(14) lnFare20	(15) lnFare80	(16) lnFare	(17) lnFare20	(18) lnFare80
GenAI	0.048***	0.032***	0.039***	0.044***	0.029***	0.028***	0.052***	0.063***	0.039***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	3,606,781	3,606,781	3,606,781	2,688,708	2,688,708	2,688,708	3,361,515	3,361,515	3,361,515

Note: Sample period is Q1 2013–Q2 2025. Each column re-estimates the baseline specification after excluding the airline named in the column title. All other specifications are the same as in Table 4. *** p<0.01, ** p<0.05, * p<0.1.

Table 11: Decomposing Airfare Effects: Quality and Capacity Channels

	(1) Markup	(2) lnFare	(3) Fare Effect	(4) lnDepartures	(5) lnSeats
AI Adoption	0.005***	0.010***		0.298***	0.294***
	(0.001)	(0.001)		(0.038)	(0.040)
$\widehat{\text{Lag HHI}}$		0.203***		24.374***	25.738***
		(0.010)		(1.418)	(1.503)
$\widehat{\text{MMC}}$		0.000***		-0.000***	-0.000***
		(0.000)		(0.000)	(0.000)
Lag Cancellation		0.001***			
		(0.000)			
ΔDelay			0.000		
			(0.000)		
ΔCancellation			0.001		
			(0.001)		
Observations	3,026,392	3,304,316	126,946	233,958	233,947
Controls	Y	Y	Y	Y	Y
Airline FE	Y	Y	Y	Y	Y
Route FE	Y	Y	Y	Y	Y
Time FE	Y	Y	N	Y	Y

Note: *Markup* is proxied by the residual from regressing fares on lagged delay and cancellation rates ($t-2$) with the full set of controls and fixed effects in Equation 3. *lnDepartures* is the log of total departures performed, and *lnSeats* is the log of total seats derived from T-100 Segment data. See Table 3 for other notes. Standard errors in parentheses are clustered by route for columns 1, 2, 4, and 5. *** p<0.01, ** p<0.05, * p<0.1.

Table 12: Impact of Enterprise AI Adoption and GenAI Platform Use on Airfare Distribution by Route Size

Enterprise AI	Small Routes			Large Routes		
	(1) lnFare	(2) lnFare20	(3) lnFare80	(4) lnFare	(5) lnFare20	(6) lnFare80
AI Adoption	-0.002	0.009***	-0.007***	0.012***	0.026***	0.012***
	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)
Observations	1,103,089	1,103,089	1,103,089	2,191,420	2,191,420	2,191,420
Controls	Y	Y	Y	Y	Y	Y
Airline FE	Y	Y	Y	Y	Y	Y
Route FE	Y	Y	Y	Y	Y	Y
Time FE	Y	Y	Y	Y	Y	Y
GenAI	Small Routes			Large Routes		
	(7) lnFare	(8) lnFare20	(9) lnFare80	(10) lnFare	(11) lnFare20	(12) lnFare80
GenAI	0.035***	0.028***	0.035***	0.061***	0.046***	0.045***
	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)
Observations	1,266,123	1,266,123	1,266,123	2,409,092	2,409,092	2,409,092
Controls	Y	Y	Y	Y	Y	Y
Airline FE	Y	Y	Y	Y	Y	Y
Route FE	Y	Y	Y	Y	Y	Y
Time FE	Y	Y	Y	Y	Y	Y

Note: Sample period is 2013Q1–2023Q3 for columns 1–6 and 2013Q1–2025Q2 for columns 7–12. For each route, we calculate the average number of tickets sold in 2013 and 2014, before any airline in our sample adopted enterprise AI. Routes above the median of this pre-period measure are classified as large routes, and those below the median are classified as small routes. See Table 3 for other notes. *** p<0.01, ** p<0.05, * p<0.1.

Table 13: Impact of GenAI Platform Use on Airfare Distribution by Baseline Scale

	(1) lnFare	(2) lnFare20	(3) lnFare80
GenAI	-0.017***	-0.043***	-0.017***
	(0.002)	(0.003)	(0.003)
GenAI × Baseline Scale	0.015***	0.018***	0.013***
	(0.000)	(0.001)	(0.000)
Observations	3,687,969	3,687,969	3,687,969
Controls	Y	Y	Y
Airline FE	Y	Y	Y
Route FE	Y	Y	Y
Time FE	Y	Y	Y

Note: Sample period is 2013Q1–2025Q2. *Baseline Scale* is defined as $\log(1 + Passenger_j^{pre})$, where $Passenger_j^{pre}$ is the number of passengers on route j before 2023Q4, when the first GenAI platform was adopted in our sample. The mean of *Baseline Scale* is 3.76, implying that the overall effect of GenAI is positive at the sample mean. All other specifications are the same as Table 4. Standard errors in parentheses are clustered by route. *** p<0.01, ** p<0.05, * p<0.1.

Online Appendix

A Timelines of Airline AI Adoption

This section elaborates on selected examples from the timeline in Figure A1, which identifies when major U.S. airlines began adopting artificial intelligence and the operational areas in which adoption first occurred. These examples illustrate the gradual diffusion of AI technologies across maintenance, scheduling, pricing, and customer-service functions since the mid-2010s. Each entry below is based on the first publicly reported or widely acknowledged instance of AI implementation drawn from reputable industry and corporate sources.

American Airlines

- **2016**: Introduced AI in baggage-handling operations to improve efficiency and reduce mishandled luggage incidents.[23]
- **2020**: Expanded AI use in flight operations and crew scheduling to improve the efficiency of disruption management.[24]
- **2021**: Adopted "Smart Gating" technology which uses AI to route aircraft to the nearest available gate.[25].
- **2022**: Partnered with Microsoft to utilize Azure as the preferred cloud platform to enhance operations with AI and machine learning.[26]

United Airlines

- **2017**: Began deploying AI-driven predictive maintenance systems through industry partnerships to enable real-time diagnostics of aircraft components and reduce unexpected failures and flight delays.[27]

[23] *CIO Dive*, Naomi Eide, "At American Airlines, AI Starts with Baggage Handling," September 28, 2018. Available at: https://www.ciodive.com/news/at-american-airlines-ai-starts-with-baggage-handling/533493/.

[24] *American Airlines Newsroom*, "Artificial Intelligence at American," available at https://news.aa.com/news/news-details/2025/Artificial-intelligence-at-American-CORP-OTH-01/default.aspx.

[25] *American Airlines Newsroom*, "Straight to the Gate This Holiday Season," November 22, 2023. Available at: https://news.aa.com/news/news-details/2023/Straight-to-the-gate-this-holiday-season-OPS-OTH-11/default.aspx.

[26] *Microsoft News Center*, "American Airlines and Microsoft Partnership Takes Flight to Create a Smoother Travel Experience for Customers and Better Technology Tools for Team Members," May 18, 2022. Available at: https://news.microsoft.com/source/2022/05/18/american-airlines-and-microsoft-partnership-takes-flight-to-create-a-smoother-travel-experience-for-cu

[27] *Predictive Aircraft Maintenance*, Mike Roemer, "Predictive Maintenance," February 2025. Available at: https://www.predictiveaircraftmaintenance.com/resources/uploads/2025/02/Mike-Roemer-final.pdf.

- **2019**: Launched the *ConnectionSaver* program, an AI-based tool that predicts whether connecting passengers can make their flights and adjusts gate holds accordingly.[28]
- **2021**: Partnered with PROS to deploy its AI-driven corporate sales platform with automated offer generation, dynamic pricing, and contract management to enhance sales efficiency and brand differentiation.[29]

Delta Air Lines

- **2015**: Implemented predictive maintenance technologies to anticipate component failures, reduce downtime, and improve operational efficiency (Ge et al., 2025b).
- **2020**: Developed an AI tool to address weather disruptions and improve flight operations.[30]
- **2023**: Began experiments with AI-driven pricing optimization in revenue-management systems.[31]

Alaska Airlines

- **2019**: Began using the Flyways AI platform to optimize air-traffic flow and improve flight efficiency.[32]

B Measuring Vendor Overlap and Platform Diffusion

As discussed in Section 3.2, we construct two route–time level indices to capture vendor commonality and platform diffusion. The Shared-Vendor Overlap Index (SVOI) measures the extent of vendor overlap among competing airlines and indicates the potential for coordinated pricing through a common vendor. The Vendor Adoption Intensity Index (VAII) measures the overall intensity of vendor use on a route, which reflects the degree of reliance on third-party pricing platforms and GenAI-enhanced pricing tools.

[28] *PR Newswire*, "United Airlines Makes Connecting the World Easier Than Ever with ConnectionSaver," June 10, 2019. Available at: https://www.prnewswire.com/news-releases/united-airlines-makes-connecting-the-world-easier-than-ever-with-connectionsaver-300864600.html.

[29] *Skift*, "PROS Gains Travel Tech Sales Uplift with United Airlines Deal," December 22, 2021. Available at: https://skift.com/2021/12/22/pros-gains-travel-tech-sales-uplift-with-united-airlines-deal/.

[30] *Aviation Today*, Woodrow Bellamy III, "Delta Develops Artificial Intelligence Tool to Address Weather Disruption, Improve Flight Operations," January 8, 2020. Available at: https://www.aviationtoday.com/2020/01/08/delta-develops-ai-tool-address-weather-disruption-improve-flight-operations/.

[31] *Aviation Week*, Christine Boynton, "Delta Air Lines Begins AI Pricing Experiments," December 6, 2023. Available at: https://aviationweek.com/air-transport/airlines-lessors/delta-air-lines-begins-ai-pricing-experiments.

[32] *Alaska Airlines Newsroom*, "Alaska Airlines and Airspace Intelligence Announce First-of-Its-Kind Partnership to Optimize Air Traffic Flow with Artificial Intelligence and Machine Learning," May 25, 2021. Available at: https://news.alaskaair.com/newsroom/alaska-airlines-and-airspace-intelligence-announce-first-of-its-kind-partnership-to-optimize-air-traff

B.1 Intuition and Examples

To build intuition for the SVOI and VAII measures, consider the following four stylized scenarios for a route j at time t:

A UA=PROS, JetBlue=PROS, DL=Fetcher (three carriers; all others=None).

B UA=PROS, JetBlue=PROS (duopoly; all others absent or None non-participants).

C DL = Fetcherr, competitor = None (duopoly with one non-vendor rival).

D All carriers = None (no vendors present on the route).

The (unweighted) SVOI for this route is given by

$$\text{SVOI}^{\mathbf{A}}_{ijt} = 1/3, \quad \text{SVOI}^{\mathbf{B}}_{ijt} = 1, \quad \text{SVOI}^{\mathbf{C}}_{ijt} = 0, \quad \text{SVOI}^{\mathbf{D}}_{ijt} = 0.^{33}$$

Ranking by SVOI (i.e., same-vendor overlap) is therefore

$$\mathbf{B} \;>\; \mathbf{A} \;>\; \mathbf{C} \;=\; \mathbf{D}.$$

Note that when pairwise weights $w_{ii't}$ are applied, the exact values of SVOI may change, but the qualitative ordering remains the same: SVOI is determined by *same-vendor* pairs rather than by the mere presence of vendors.

On the other hand, the VAII for this route is given by

$$\text{VAII}^{\mathbf{A}}_{jt} = 3/3 = 1, \quad \text{VAII}^{\mathbf{B}}_{jt} = 2/2 = 1, \quad \text{VAII}^{\mathbf{C}}_{jt} = 1/2 = 0.5, \quad \text{VAII}^{\mathbf{D}}_{jt} = 0.$$

Ranking the cases by VAII (i.e., how algorithmic the route is) yields:

$$\mathbf{A} \;=\; \mathbf{B} \;>\; \mathbf{C} \;>\; \mathbf{D}.$$

B.2 Computation Steps

We now outline the procedure for computing the SVOI and VAII measures. For each year-quarter t:

1. Construct the airline-time vendor map $i \mapsto v_{it}$ (vendor assignment for each airline i at time t).
2. For each route j, define the set of serving airlines $\mathcal{J}_{jt}$ and their market shares s_{ijt}.
3. Compute airline network overlap for the baseline period. Specifically, measure the fraction of routes jointly served by airlines i and i' relative to the total distinct routes served by either carrier (a Jaccard-style index); denote this $\Omega_{ii'} \in [0, 1]$.

[33] In case A, with three airlines competing ($|\mathcal{J}_{jt}| = 3$), there are three possible pairs; only UA and JB share a vendor, so exactly one of three pairs matches, giving SVOI = 1/3. In case B, two airlines compete ($|\mathcal{J}_{jt}| = 2$) and form a single pair sharing the same vendor, so SVOI = 1. In cases C and D, no two airlines share the same vendor, hence SVOI = 0.

4. Set pairwise weights $w_{ii't} = \Omega_{ii'}$.

5. Compute SVOI_{ijt} using Equation (1), and compute VAII_{jt} using Equation (2).

6. Merge the computed indices SVOI_{ijt} and VAII_{jt} back to airline-route-time cells (i, j, t), and estimate Equation (4).

Corner cases: if $|\mathcal{J}_{jt}| < 2$, set $\text{SVOI}_{ijt} = 0$; if $|\mathcal{J}_{jt}| = 0$, set both indices to 0.

C Additional Robustness Checks

C.1 Measurement Error in Adoption Timing

There may be concerns about potential measurement error in treatment timing, as AI adoption dates from public sources could involve phase-in periods or partial rollout. To address this, we estimate event-time bin specifications that exclude the transition window, where we drop observations one or two quarters before and after the recorded adoption date. The estimated coefficients, as reported in Table A2, remain qualitatively similar. Thus, even after omitting the transition period, the estimated treatment effect is stable, indicating that adoption timing is measured with reasonable accuracy and not driven by anticipation or gradual phase-in effects.

C.2 COVID-19 Disruptions

One might be concerned that routes where AI adoption is most prevalent, typically larger markets, also experienced more severe and heterogeneous disruptions during the COVID-19 pandemic. While time fixed effects absorb the common pandemic shock, they do not capture route-specific variation. To address this, we follow Ge et al. (2025a) and re-estimate our baseline enterprise AI adoption specifications with additional controls for state-level COVID policies at each route endpoint, including phases of travel restrictions and vaccine rollouts. The estimates, as reported in Table A3, remain robust.

D Theoretical Framework for GenAI Pricing and Vendor Adoption

In this section, we develop a simple theoretical framework for GenAI pricing and vendor adoption. The model builds on the partial profit internalization framework commonly used in the industrial organization literature (Bresnahan 1982; Bernheim and Whinston 1990). In this framework, firms may place a positive weight on rivals' profits, which can capture coordinated conduct. Recent work has applied similar objective functions in the airline industry; for example, Azar and Ribeiro (2021) models managers as maximizing a weighted sum of own and rival firms' profits reflecting common ownership.

D.1 Setting

Demand and profit function with vendor adoption intensity On a route j, there exist $N \geq 2$ airlines $i \in \{1, 2, \ldots, N\}$ airlines. Each airline chooses a price $p_i \in \mathcal{P} = [\underline{p}, \bar{p}] \subset \mathbb{R}$ and let $\mathbf{p} = (p_1, p_2, \ldots, p_N)$.

Each airline can adopt a GenAI pricing vendor $v(i) \in \{0, 1, \ldots, V\}$ where $v(i) = 0$ means "no vendor adoption." Let $m \in [0, 1]$ denote the route-level vendor adoption intensity which equals the fraction of airlines on the route that use any GenAI pricing vendor.

The demand function for each airline i is denoted by $q_i(p, m)$ where q_i is a function as

$$q_i : \mathcal{P}^N \times [0, 1] \to \mathbb{R}_+.$$

Marginal costs are constant for each airline, $c_i \in (\underline{p}, \bar{p})$. The profit function can be written as

$$\pi_i(\mathbf{p}; m) = (p_i - c_i) q_i(\mathbf{p}; m).$$

Same-vendor overlap We define the set of rivals who adopted the same vendor as airline i as S_i,

$$S_i \equiv \{\, i' \neq i : v(i') = v(i) \neq 0 \,\}.$$

Each airline can partially internalize rivals' profits when they adopt a common vendor through coordination, so the profit function can be expressed as

$$\tilde{\pi}_i(\mathbf{p}; m) \equiv \pi_i(\mathbf{p}; m) + \theta \sum_{i' \in S_i} \pi_{i'}(\mathbf{p}; m).$$

We refer to the profit function above as effective profit function. θ represents the intensity of softened competition via adopting the same vendor. When $\theta = 0$, airlines use a standard Bertrand pricing, and $\theta \in (0, 1)$ shows the case where coordinated conduct is present and airlines take the profits of other airlines with the same vendor into consideration when pricing.

Given $(m, v(\cdot), \theta)$, airlines simultaneously choose their price p_i to maximize $\tilde{\pi}_i(\mathbf{p}, m)$.

Assumptions We provide assumptions of the model before we discuss the implications of the model on adopting and sharing vendors. Assumption 1 is needed for a well-behaving best response mapping, so each profit maximization problem has a solution.

Assumption 1 (Best-response function) *For each firm i, the demand function $q_i(p; m)$ is twice continuously differentiable on $\mathcal{P}^N$ for each $m \in [0, 1]$. The effective profit function $\tilde{\pi}_i(p; m)$ is strictly concave in p_i for every (p_{-i}, m). Therefore, the best response function*

$$BR_i(p_{-i}; m) = \arg\max_{p_i \in \mathcal{P}} \tilde{\pi}_i(p_i, p_{-i}; m)$$

is single-valued.

Assumption 2 states that prices are starategic complements. This property is standard in differentiated Bertrand competition models, where an increase in a rival's price raises

the marginal profitability of increasing one's own price. In airline markets, products are differentiated by departure times, service quality, and frequent flyer programs, which leads to strategic complementarity in prices across competing carriers.

Assumption 2 (Strategic complements) *For all $i \neq i'$,*

$$\frac{\partial^2 \tilde{\pi}_i(\mathbf{p}; m)}{\partial p_i \partial p_{i'}} \geq 0 \quad \textit{for all } (\mathbf{p}, m) \in \mathcal{P}^N \times [0, 1].$$

We also introduce an assumption that captures the idea that broader adoption of GenAI pricing vendors intensifies price competition in the market. As more airlines adopt vendors, prices are adjusted more frequently and it becomes easier to observe competitors' pricing decisions, which leads to less pricing frictions. As a result, the marginal effect of profits increasing from raising prices declines as more airlines adopt GenAI pricing vendors.

Assumption 3 (Vendor adoption and competition) *For each airline i,*

$$\frac{\partial^2 \tilde{\pi}_i(\mathbf{p}; m)}{\partial p_i \partial m} \leq 0 \quad \textit{for all } (p, m) \in \mathcal{P}^N \times [0, 1].$$

Equilibrium A pure-strategy Bertrand Nash equilibrium is a price vector $\mathbf{p}^*(m, v, \theta) \in \mathcal{P}^N$ such that for each i,

$$p_i^* = \arg\max_{p_i \in \mathcal{P}} \tilde{\pi}_i(p_i, p_{-i}^*; m).$$

Under Assumptions 1 and 2, the airlines' pricing game is supermodular (Topkis 1998; Milgrom and Shannon 1994), and hence the set of equilibria is nonempty and forms a complete lattice. In particular, there exist a smallest and a largest equilibrium price vectors component-wise. We focus on the largest equilibrium price vector, denoted by $\mathbf{p}^{\max}(m, v, \theta)$.

D.2 Implications

We derive two propositions from the model developed above. The first concerns how same-vendor overlap affects equilibrium prices through coordination. In particular, when more airlines share the same vendor, equilibrium prices weakly increase. Formally,

Proposition 1 (Same-vendor adoption and increasing prices) *For two vendor adoption assignments v and v' where $S_i(v) \subseteq S_i(v')$,*

$$\mathbf{p}^{\max}(m, v', \theta) \geq \mathbf{p}^{\max}(m, v, \theta),$$

where prices are compared component-wise.

Increasing the shared vendor set $S_i(v)$ increases the partially internalized profit term, $\theta \sum_{j \in S_i} \pi_j$ in each airline's profit function. If airline i cuts price p_i, then its own price $\pi_i(\mathbf{p}; m)$ would increase but rivals' profits fall, hence $\sum_{j \in S_i} \pi_j$ falls, reducing the incentive to undercut prices. Specifically, expanding S_i increases the marginal payoff of increasing

price p_i in $\tilde{\pi}_i(\mathbf{p}; m)$. Under Assumption 2, increasing airline i's prices weakly increases other airlines' prices which leads to higher largest equilibrium. This implies that the shared vendor acts as a coordination device from the partially internalized profits among its clients.

Next, Assumption 3 implies that the effective profit function $\tilde{\pi}_i$ exhibits decreasing differences in (p_i, m). Hence, an increase in vendor adoption intensity reduces the marginal return to raising price. By the monotone comparative statics results of Topkis (1998) and Milgrom and Shannon (1994), the firm's optimal price is non-increasing in m. Because prices are strategic complements (Assumption 2), the largest equilibrium price vector is also weakly decreasing in m. Formally,

Proposition 2 (Vendor adoption intensity and decreasing prices) *Given vendor assignment v and θ, under Assumptions 1–3, the largest equilibrium price is weakly decreasing in m. For $m' \geq m$,*

$$\mathbf{p}^{\max}(m', v, \theta) \leq \mathbf{p}^{\max}(m, v, \theta),$$

where prices are compared component-wise.

Since m represents the vendor adoption intensity in a route, Proposition 2 predicts that as more airlines adopt GenAI vendors, equilibrium prices weakly decline.

Figure A1: AI Adoption Timelines by Airline

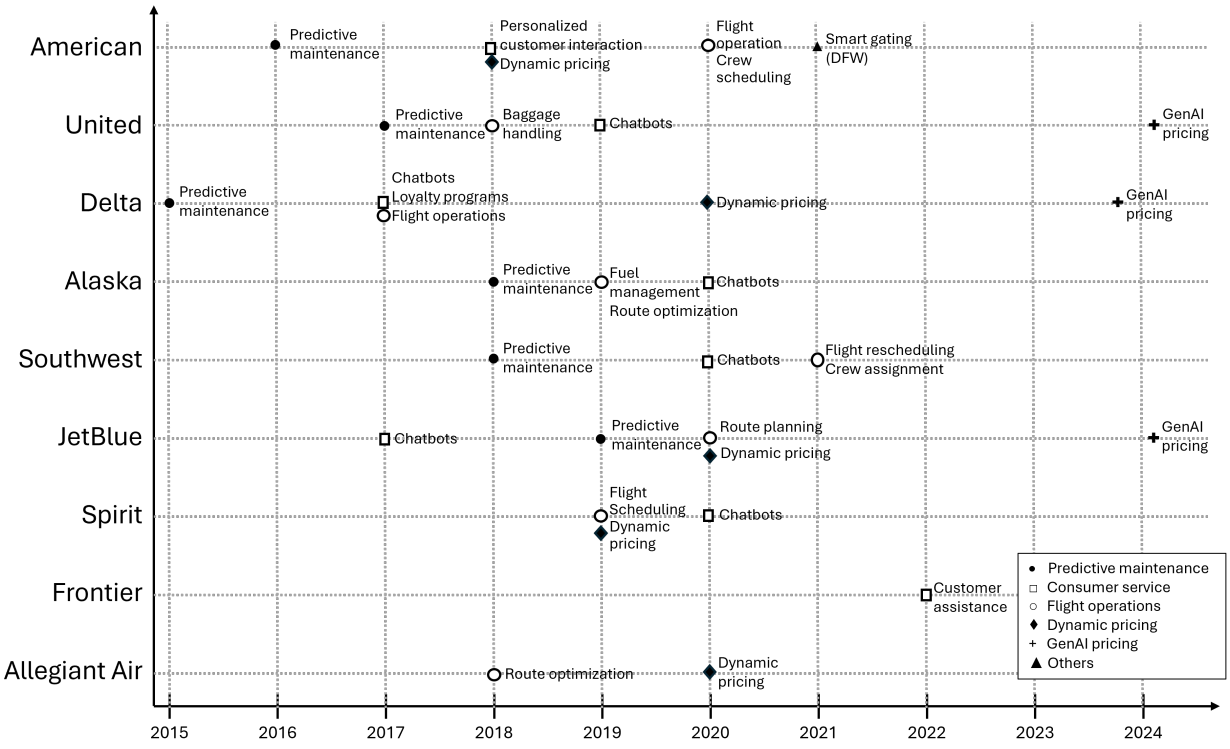
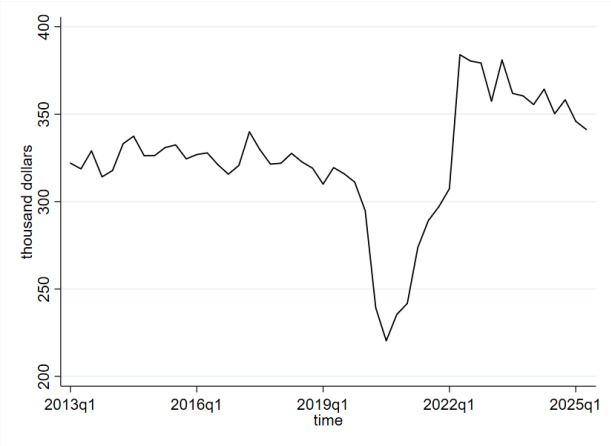
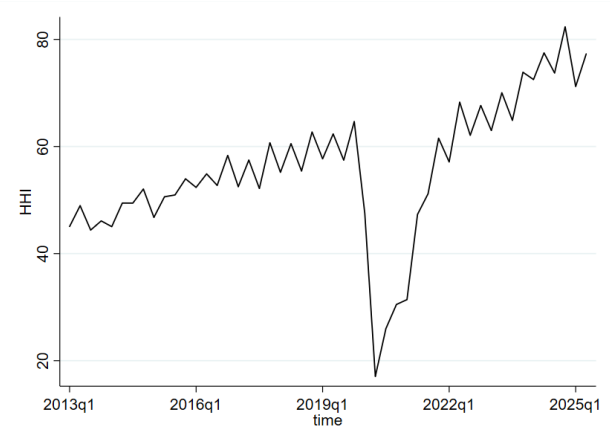


Figure A2: Evolution of Airfares, On-Time Performance, Market Concentration, and Air Quality

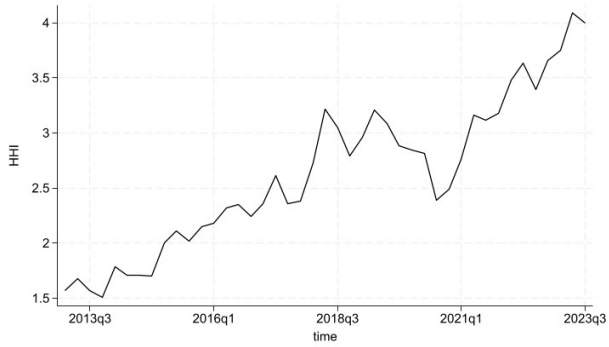
(a) Airfares



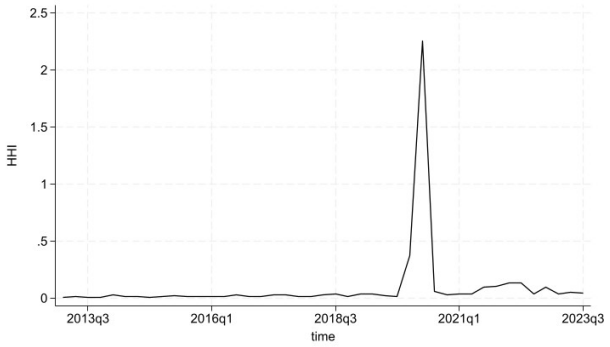
(b) HHI



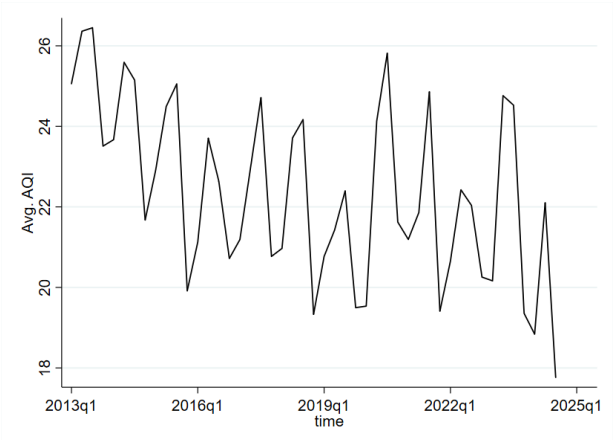
(c) Delay Time



(d) Cancellation (%)



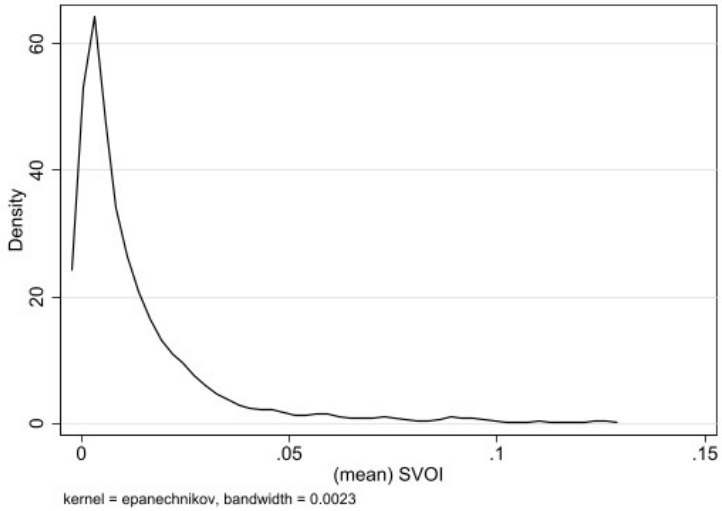
(e) AQI



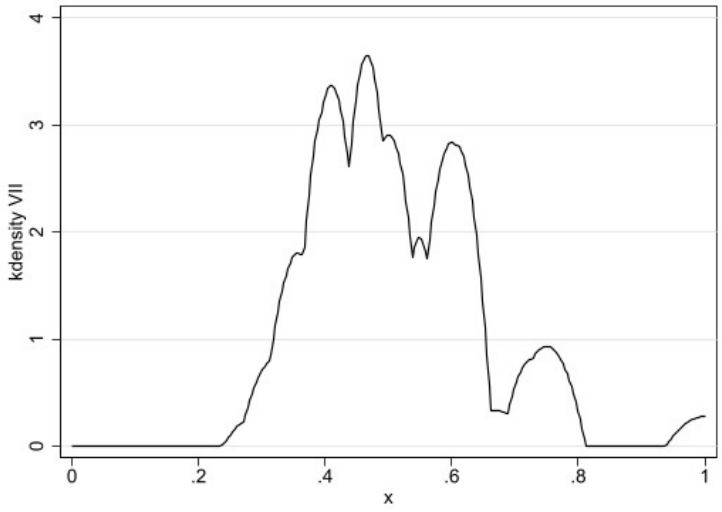
Note: Delay time and cancellation rate exclude delays due to weather related reasons.

Figure A3: Distribution of GenAI Indices

(a) SVOI



(b) VAII



Note: The figure only includes data from 2025Q2 and observations with positive index values. Each histogram is constructed using the average SVOI and VAII values at the route level.

Table A1: Impact of Enterprise AI Adoption (First Stage)

	(1)
Lag lnFare	27.691***
	(9.034)
Lag Admin. Expense	0.049***
	(0.018)
Lag Res. Expense	0.228***
	(0.076)
Lag Line Svc. Expense	-0.105***
	(0.034)
Lag #Routes	-686.707
	(450.943)
Lag Distance	-26.874***
	(8.446)
Lag #Tickets	-0.000***
	(0.000)
Wald χ^2	313.92
Observations	247

Note: This table reports the first-stage estimates for Tables 2 and 3. The specification is a probit model with year fixed effects, where the dependent variable equals 1 if an airline adopts enterprise AI and 0 otherwise, using airline-time level data. *Tickets* denotes the total number of tickets sold by an airline in each year-quarter. *** p<0.01, ** p<0.05, * p<0.1.

Table A2: Robustness Check: Impact of Enterprise AI Adoption on Airfare Distribution (Event Time Bins)

	(1) lnFare	(2) lnFare	(3) lnFare	(4) lnFare
AI Adoption	0.055***	0.063***		
	(0.007)	(0.011)		
GenAI			0.058***	0.022***
			(0.010)	(0.004)
Observations	834,705	756,466	3,363,175	3,320,037
Controls	Y	Y	Y	Y
Airline FE	Y	Y	Y	Y
Route FE	Y	Y	Y	Y
Time FE	Y	Y	Y	Y

Note: Sample period is 2013Q1–2023Q3 for columns 1–2 and 2013Q1–2025Q2 for columns 3–4. Column 1 excludes 1 quarter before and after AI adoption, and Column 2 excludes 2 quarters before and after AI adoption. Column 3 excludes 1 quarter before and after GenAI pricing adoption, and Column 4 excludes 2 quarters before and after GenAI pricing adoption. See Table 3 for other notes. *** p<0.01, ** p<0.05, * p<0.1.

Table A3: Robustness Check: Impact of Enterprise AI Adoption on On-Time Performance and Air Quality (Controlling for COVID-19 Policies)

	(1) Delay Time	(2) Cancellation (%)	(3) AQI
AI Adoption	-0.911***	0.040***	-2.190***
	(0.061)	(0.004)	(0.049)
Stay-at-Home Started	0.082	0.009*	-0.082
	(0.057)	(0.005)	(0.119)
Stay-at-Home Ended	-0.181***	-0.025***	-0.288**
	(0.063)	(0.005)	(0.139)
Quarantine Started	0.105*	-0.005	0.189
	(0.059)	(0.007)	(0.121)
Quarantine Ended	-0.003	-0.009	-0.697***
	(0.062)	(0.006)	(0.130)
Vaccine	-0.002	0.008**	0.554***
	(0.057)	(0.004)	(0.129)
Observations	3,906,338	3,906,338	2,922,924
Controls	Y	Y	Y
Airline FE	Y	Y	Y
Route FE	Y	Y	Y
Time FE	Y	Y	Y

Note: Sample period is 2013Q1–2023Q3. *Stay-at-Home* and *Quarantine* are dummy variables that equal 1 if the origin or destination state has implemented a stay-at-home order and quarantine requirement for out-of-state travels, respectively, and zero if the travel restrictions have not been enacted or have been lifted. *Vaccine* is a similarly defined dummy that equals 1 if COVID-19 vaccines (by Moderna, Pfizer, or Johnson & Johnson) have become available to the general public in either the origin or destination state. *AQI* is the average AQI weighted by fuel expenses multiplied by passenger shares. See Table 3 for other notes. *** p<0.01, ** p<0.05, * p<0.1.

Table A4: Robustness Check: Impact of GenAI Platform Use and Vendor Sharing on Airfare Distribution (First Stage for Control Function Approach)

	(1) SVOI	(2) VAII
Lag lnFare	0.000***	-0.005***
	(0.000)	(0.000)
Lag Admin. Expense	-0.362***	-22.895***
	(0.033)	(1.096)
Lag Res. Expense	1.426***	79.719***
	(0.137)	(2.931)
Lag Line Svc. Expense	-0.697***	-145.416***
	(0.086)	(3.428)
Lag Distance	0.000	-0.001***
	(0.000)	(0.000)
Lag BLP GenAI	0.001***	0.104***
	(0.000)	(0.001)
Lag Origin SVOI	1.430***	
	(0.097)	
Lag Dest. SVOI	1.371***	
Observations	4,684,374	4,684,374
F-stat	131.69	1255.80
Controls	Y	Y
Airline FE	Y	Y
Route FE	Y	Y
Time FE	Y	Y

Note: This table reports the first-stage estimates for Table 6. BLP SVOI is the sum of competitors' SVOI in the same market, and BLP GenAI is the number of other airlines that adopted GenAI in the same market. Administration expenses, Reservation and sales expenses, and line servicing expenses are weighted by share of ticket numbers. Origin and destination SVOI are the average of SVOI The sample includes Delta, United, American, Frontier, Southwest, and JetBlue. Standard errors in parentheses are clustered by route. *** p<0.01, ** p<0.05, * p<0.1.