Artificial Intelligence, Operational Efficiency, and Price Coordination*

Qi Ge† Minhae Kim‡ Myongjin Kim§

Abstract: We study how artificial intelligence adoption affects pricing, service quality, environmental externalities, and consumer welfare in the U.S. airline industry. Enterprise AI adoption improves on-time performance and local air quality while increasing average fares and compressing price dispersion. Generative AI-enhanced pricing platforms produce larger fare effects whose direction depends on carrier adoption patterns. Shared vendor relationships are associated with higher fares, especially at the lower end of the fare distribution, whereas broader platform diffusion is associated with lower fares. Welfare calculations indicate net gains for consumers. Overall, our findings suggest a key trade-off between efficiency improvements and tacit coordination risks.

1 Introduction

Artificial intelligence (AI) uses data inputs to generate predictions, recommendations, or decisions (McElheran et al., 2024), and AI systems are transforming firm decision-making across industries, ranging from manufacturing and logistics to healthcare and transportation. Recent advances in machine learning methods, computational power, and data availability have made AI especially valuable in complex, data-intensive environments, where firms must continuously process information on demand, costs, capacity, and competitors (Acemoglu et al., 2022; Babina et al., 2024). As AI systems become increasingly embedded in operational and business decisions, their impacts extend beyond firm productivity to pricing behavior, market power, and consumer welfare.

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. The global aviation AI market was estimated at $1.01 billion in 2024 and is projected to reach $32.50 billion by 2033.[1] 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

*The authors received no financial support for this research and have no financial relationships or other potential conflicts of interest to disclose.

†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]Source: Pavan Warade, "Artificial Intelligence in Aviation Market Size, Share & Trends Analysis Report," *Straits Research*, June 2025, `https://straitsresearch.com/report/artificial-intelligence-in-aviation-market`.

enhance predictive maintenance, flight operations and scheduling, baggage handling, and customer service.[2] Beyond operations, AI is playing an increasingly integral role in airline competition and pricing. Enterprise AI systems can analyze large volumes of data to support reactive price optimization in response to demand fluctuations, booking patterns, seat availability, and competitor prices. More recently, generative AI (GenAI) applications have added a simulation and recommendation layer to these capabilities, marking a shift in the pricing paradigm from primarily forecast-based fare optimization toward pricing recommendations based on product attributes, market conditions, and competitive dynamics. For example, Delta Air Lines' recent GenAI pricing pilot, referred to as a "super analyst," drew public scrutiny over whether AI pricing could lead to surveillance pricing or individualized fare setting, despite Delta's later statement rejecting such claims.[3]

Even without relying on passenger-specific data, the use of AI-based pricing algorithms raises significant antitrust concerns. Algorithmic systems can respond to rivals in ways that mimic tacit coordination when repeated interactions generate feedback loops in pricing. Airlines that adopt similar AI solutions from shared vendors may also unintentionally coordinate pricing through common algorithms and data inputs (Aleksenko and Miklós-Thal, 2026). These concerns call for a comprehensive analysis of impacts of AI technologies and their implications for consumer welfare.

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 analyze AI adoption as a carrier-level technology shock with heterogeneous effects across route markets and 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 between 2013Q1 and 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, we estimate how carrier AI adoption affects route-level outcomes by comparing airline-route pairs before and after adoption to those of non-adopting carriers. Our empirical framework does not require observing route-level AI deployment. Instead, we estimate route-level exposure to carrier-level AI shocks, analogous to airline merger retrospectives that identify product market effects using heterogeneous exposure to airline-level events. We find that enterprise AI

[2]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.

[3]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`. See also Letter from Senators Ruben Gallego, Richard Blumenthal, and Mark R. Warner to Ed Bastian, CEO of Delta Air Lines, July 21, 2025, which raised concerns about Delta's planned expansion of AI-based revenue management and individualized fare setting. Delta stated that its AI pricing functionality is an analyst-overseen decision support tool that uses aggregated data, including purchasing data and demand forecasts for specific routes and flights, and does not use personal information to set individualized fares. See 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`.

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 is larger at the bottom of the fare distribution than at the top, which compresses price dispersion and creates greater uniformity across routes. To further assess the conduct channel, we test whether enterprise AI adoption raises markups above the Bertrand–Nash benchmark. The estimated conduct index is positive and implies a markup increase of about 12% above the Bertrand benchmark, which is consistent with the 15% margin increase documented by Assad et al. (2024) in the German retail gasoline market. These results suggest that the fare increases associated with enterprise AI adoption reflect changes in pricing conduct and market power, not just shifts in demand composition or fare adjustments based on standard pricing inputs.

To explore potential price coordination channels, we focus on the use of GenAI-enhanced pricing platforms add a market-simulation and recommendation layer to traditional revenue management. Rather than simply improving demand forecasts within an existing fare optimization system, these platforms combine internal airline data with external market signals to infer demand states, simulate competitive scenarios, and generate analyst-facing pricing recommendations.In airline pricing, adoption of these platforms occurs through carrier-level vendor relationships instead of route-specific deployment decisions. Their pricing effects may therefore emerge gradually over time and across route markets as airlines integrate the platforms into existing revenue management workflows. We exploit variation in vendor adoption across airlines, as well as common-vendor overlap and platform diffusion across route markets, to examine how potential coordination may arise when carriers rely on a common algorithm vendor and as these platforms become more widespread. We find that the pricing impacts become even larger for GenAI-enhanced pricing platforms, with fares rising throughout 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 on a given route affects equilibrium fares through common vendor overlap and platform diffusion. In this setting, greater same-vendor overlap dampens competition by expanding the set of rivals whose profits are partially internalized, which in turn raises equilibrium prices. Meanwhile, the model predicts that as platform adoption intensity increases within a route, the marginal return to raising prices declines, which puts downward pressure on equilibrium fares. These interpretations also align with related work showing that greater overlap in common outsourcing relationships can soften competition and raise fares (Aryal et al., 2024), as well as the recent literature 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 the enterprise AI and GenAI analyses. Because our empirical design treats AI adoption as a carrier-level technology shock with heterogeneous effects across route markets, similar to retrospective merger studies, the baseline estimates may be affected by staggered treatment timing, alternative measures of AI adoption, and unobserved time-varying carrier-level or route-specific shocks that coincide with adoption. We address these concerns using staggered-adoption estimators, alternative adoption measures, event-time bin analysis, richer fixed effects and interactive-factor specifications, placebo tests, and sensitivity checks that assess whether the results are driven by any single airline. Across these exercises, our main results remain 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 product market impacts, ranging from on-time performance and air quality to fares and their distributions. Our findings provide new evidence on the trade-off between efficiency gains and coordination risks in the adoption of AI technologies.

Next, to our knowledge, this paper is also among the first to examine the economic impacts of exposure to *generative* AI technologies and distinguish these effects from those associated with enterprise AI. Recent work emphasizes that the economic effects of GenAI depend on complementary assets and control over key inputs, as well as how GenAI products interact with non-GenAI alternatives (Azoulay et al., 2025; Goldberg and Lam, 2025). Our analysis provides novel empirical evidence from the airline industry that the use of pricing platforms with generative components, introduced as an additional layer integrated into the existing enterprise pricing systems, can raise fares in settings where rivals share the same vendor, whereas greater diffusion within 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 behavior 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 depends on the informational environment in which firms interact. Our findings are consistent with this view and complement recent work on third-party pricing algorithms, which suggests that common pricing systems can affect competitive outcomes when they process information from competing firms (Aleksenko and Miklós-Thal, 2026).

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 affect competitive outcomes across route markets. 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 have important welfare and policy implications. Consumer welfare calculations based on our demand estimates indicate that enterprise AI adoption generates welfare gains from improved service quality of about $888.8 million, losses from higher fares of about $604.5 million, and environmental benefits of roughly $424.9 million, yielding a substantial net positive effect for consumers. We also estimate that the fare increases associated with GenAI platform use result in additional consumer welfare losses of about $199.5 million. On balance, these estimates suggest that AI adoption can enhance overall consumer welfare by improving product quality and reducing negative environmental externalities despite modest price increases. At the same time, our findings underscore the need for careful regulatory oversight. The risks of algorithmic coordination depend not only on the type of AI pricing tool, but also on how it is deployed and the competitive environment (Miklós-Thal and Tucker, 2024). As AI-driven pricing becomes more widespread, regulators should pay close attention to vendor relationships, market structure, and emerging concerns over personalized pricing.

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 and gradually incorporated machine learning tools for demand forecasting and fare optimization in the 2000s (Ge et al., 2025b). Since the mid-2010s, airlines have increasingly adopted enterprise AI systems across maintenance, operations, customer service, and pricing. Unlike standalone machine-learning models designed for specific prediction tasks, enterprise AI integrates forecasting and optimization algorithms into broader organizational processes using operational, booking, inventory, and pricing data. In revenue management, these systems support forecasting, inventory control, and fare optimization within existing pricing architectures.

Early applications focused on predictive maintenance, allowing airlines to anticipate equipment failures and reduce downtime (Ge et al., 2025b). Adoption later expanded to flight operations, crew scheduling, customer service, and pricing.[4] More recently, airline pricing has begun to shift from traditional enterprise revenue-management systems toward GenAI-enhanced pricing platforms supplied by vendors such as PROS, Fetcherr, and FLYR (Ge et al., 2025b). These platforms are discussed separately below because they introduce a distinct market-simulation and recommendation layer into the pricing process.

Generative AI has also recently been introduced into airline pricing platforms. GenAI produces new content, such as text, code, scenarios, or explanations, based on large language models (LLMs) trained on broad corpora, and can generate summaries and recommendations from pricing-related inputs and outputs. In airline pricing, these GenAI-enhanced systems leverage internal airline data

[4]See Geske et al. (2024) for a review.

and external market information, such as historical sales, booking and demand patterns, competitor prices, operational capacity, and macroeconomic indicators, to model product- and market-level demand conditions rather than passenger-specific willingness to pay.[5] This difference also distinguishes GenAI pricing from personalized pricing, as public statements by airlines and vendors consistently frame GenAI systems as learning from route-, flight-, and inventory-level demand conditions rather than passenger-specific characteristics.[6] Moreover, GenAI pricing platforms differ from enterprise AI in where they enter the pricing process. Traditional enterprise AI in revenue management forecasts demand, estimates elasticity, manages seat inventory, and optimizes fares within an existing RM architecture. GenAI pricing platforms, on the other hand, sit closer to the market-simulation and recommendation layer, where they use internal airline data and external market conditions to simulate market states and provide analyst-facing pricing recommendations.

Given carriers' different GenAI adoption paths, the diffusion of GenAI in airline pricing is best understood as a shift in the pricing paradigm instead of a wholesale replacement of earlier systems. Conceptually, GenAI adds a new decision layer that moves pricing from forecast-and-optimize revenue management toward market simulation and recommendation, while remaining integrated into existing airline systems and vendor platforms. For carriers such as United and JetBlue, which use PROS as their GenAI pricing vendor, the GenAI layer builds on an incumbent revenue management and pricing platform.[7] For Delta, the GenAI layer entered through a new partnership with Fetcherr, a vendor focused on GenAI-based pricing and inventory management for airlines. Other carriers, such as American, continued to rely on in-house systems or legacy pricing tools and therefore did not necessarily receive a comparable vendor-supplied GenAI layer.

Where adopted, GenAI functions as an additional decision support layer that translates product- and market-level information into analyst-facing recommendations. On the other hand, in settings where carriers rely on the same algorithm vendor with GenAI capacity, this added layer may make airlines more likely to process the same market signals and adjust their pricing responses in similar ways.

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 a workflow layer,[8]

[5]For example, Google Cloud describes Fetcherr's Large Market Model as processing internal and external data, including historical sales, real-time demand, competitive prices, macroeconomic indicators, and operational capacity. See Google Cloud, "Fetcherr Case Study," `https://cloud.google.com/customers/fetcherr-ai`. Reuters also reports that Google Cloud partnered with Air France-KLM to deploy generative AI technology on the airline group's data, while Air France-KLM would retain full ownership and control of its data. See Reuters, "Google Cloud partners with Air France-KLM on AI technology," December 4, 2024, `https://www.reuters.com/technology/artificial-intelligence/google-cloud-partners-with-air-france-klm-ai-technology-2024-12-04/`.

[6]For example, Delta describes its AI pricing functionality as an analyst-overseen decision support tool that uses aggregated data, including purchasing data and demand forecasts for specific routes and flights, and states that it does not use personal information to set individualized fares. See 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`.

[7]PROS describes its airline revenue management products as supporting dynamic pricing, network optimization, and analyst-facing recommendations, and its Revenue Management Agents blend predictive AI with LLM technology to provide autonomous decision support. See `https://pros.com/ai-agents`.

[8]See Wang et al. (2024a) for an overview of agentic AI systems.

although 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 about 20 percent by year-end.[9]

Figure A1 in the E-Companion 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.

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.[10] 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 facilitate coordinated pricing even without any generative AI component. First, enterprise AI can strengthen coordination through a fast detection, fast response mechanism in markets with repeated interaction and transparent rival prices. Airlines observe many of the same market signals, including rivals' fares, booking pace, schedules, and inventory conditions. Pricing and inventory modules can process these signals quickly and repeatedly, so that one airline's fare adjustment is detected and matched more rapidly by rival pricing systems. Over time, such responsiveness can reduce the gains from aggressive undercutting, reinforce parallel price movements, and generate pricing patterns that resemble tacit coordination. Recent experimental evidence supports this prediction. For example, in a simulated airline market, 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 generally, 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 may enhance efficiency, it can also facilitate coordination by making rivals' fare adjustments easier to detect and

[9] 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`.

[10] 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.

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).

Furthermore, enterprise AI adoption can affect pricing through cost-side channels. Platforms can lower operating costs by automating routine tasks and improving forecasting and 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, which 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, namely algorithmic responsiveness associated with platform diffusion and coordination risk arising from shared algorithm vendors. The distinction between enterprise AI and GenAI becomes particularly relevant for the coordination channels discussed below because GenAI does not simply make existing optimization faster; it can also affect how market signals are translated into pricing recommendations.

First, when rival carriers share the same GenAI-enhanced pricing platform, their pricing recommendations and responses are more likely to converge, which can narrow the scope for independent price setting and raise coordination risk. Note that the relevant mechanism here does not require direct data sharing among airlines. When rival carriers use the same GenAI-enhanced pricing vendor, the vendor may supply a similar market representation, scenario engine, objective structure, update process, and mapping from market signals into pricing recommendations. As a result, competing carriers may receive pricing recommendations generated from similar policy rules, which can make their pricing responses more aligned without direct access to one another's proprietary data. The shared-vendor mechanism is analogous to Aryal et al. (2024), who find that major carriers subcontracting to the same regional airline exhibit higher average airfares. On the other hand, 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). The net effect of shared GenAI-enhanced platforms is thus an empirical question depending on how

vendor incentives and algorithmic design affect pricing responses.

Second, GenAI-enhanced pricing platforms can also affect pricing as they become more widespread among carriers competing on the same route, even when rivals rely on different vendors. While GenAI is not itself a learning algorithm, controlled experiments with Q-learning algorithms in Calvano et al. (2020) provide useful intuition for how faster and more automated pricing responses can stabilize pricing outcomes. This platform-diffusion channel differs from the same-vendor mechanism because it does not require rival airlines to use the same market model, recommendation engine, or vendor architecture. Instead, it captures route-level diffusion of GenAI-enhanced pricing platforms, as more carriers use algorithmic systems to process similar market signals and update pricing recommendations more frequently.

Theoretically, broader platform diffusion can cut in either direction. On one hand, as GenAI-enhanced pricing platforms become more common on a route, pricing may become more systematic and predictable. Because vendor systems are often built around similar optimization principles, they may respond in parallel to common market signals, reducing uncertainty about rivals' pricing behavior and making aggressive undercutting less likely (Wang et al., 2024b). Such stability can facilitate parallel pricing behavior and result in lower fare dispersion. On the other hand, broader platform diffusion can also sharpen competitive discipline. When many carriers use AI pricing tools, high-price deviations may become less sustainable because rivals can identify and respond to profitable pricing opportunities more quickly. In this case, broader platform diffusion can place downward pressure on fares by improving the speed and accuracy of demand forecasting and price adjustments across the market. As with the shared-vendor channel, the net price effect of platform diffusion is therefore a priori indeterminate. Our estimates suggest that the competitive-discipline force dominates: market-wide vendor adoption is associated with lower fares. Table A1 in the E-Companion summarizes the conceptual distinctions, pricing channels, and empirical measures.

3 Data

3.1 Data Sources and Sample Construction

We construct an airline–route–quarter panel of U.S. domestic airline operations from 2013Q1 to 2025Q2. 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.[11]

Our primary measures for an airline's quality provisions focus on its on-time performance record, which is 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 (OTP) 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 non-

[11] See Section A in E-companion for examples of adoption timelines and sources.

directional airport pairs. 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 percentage of flights canceled for reasons other than weather. By excluding weather-related delays and cancellations, these measures capture operational disruptions driven by airline performance as opposed to 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 price distribution, which are respectively used to capture lower- and upper-tail fare segments, often associated with leisure and business demand.

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 of the counties where the origin and destination airports are located.

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 competing 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.[12] 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.[13]

More formally, in our carrier-route panel, each airline i serving a non-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.

[12]Note that algorithm vendors, such as PROS and Fetcherr, license their pricing platforms to airlines on an individual-client basis, and 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 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.

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

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:

$$\text{SVOI}_{ijt} = \frac{1}{W_{ijt}} \sum_{i' \in \mathcal{J}_{jt}, i' \neq i} w_{ii't} \mathbb{1}\{v_{it} = v_{i't} \neq \text{NONE}\}, \tag{1}$$

where $\mathcal{J}_{jt}$ is the set of airlines serving route j at time t and

$$W_{ijt} = \sum_{i' \in \mathcal{J}_{jt}, i' \neq i} w_{ii't}.$$

By construction, $\text{SVOI}_{ijt} \in [0, 1]$, and the measure captures, for focal airline i on route j at time t, the weighted share of its route competitors that rely on the same third-party pricing 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 $\text{SVOI}_{ijt} \equiv 0$.

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):

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

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. Section B provides further examples and computational steps for the SVOI and VAII measures, and Section 6.3 discusses potential endogeneity concerns related to these indices.

3.3 Descriptive Patterns

Table A2 and Figure 1 in E-companion report summary statistics and descriptive patterns by enterprise AI adoption and GenAI platform use. Overall, AI-adopting routes exhibit shorter delays, higher fares, and greater fare dispersion, while routes with greater GenAI platform penetration display higher average fares and greater vendor overlap. More detailed descriptive evidence is discussed in E-Companion Section C.

4 Empirical Specifications

To identify the causal impact of AI adoption, we employ a difference-in-differences design that compares airline-route pairs of adopting carriers with those of non-adopting carriers across the pre- and post-adoption periods. Our empirical framework is analogous to airline merger retrospectives. The relevant shock occurs at the carrier or vendor-platform level, while product market effects are estimated using heterogeneous exposure across route markets. In merger studies (e.g., Prince and Simon 2017), researchers observe the carrier-level merger event and estimate market-level effects by comparing routes that are differentially exposed to the merging carriers. Similarly, we do not observe the exact route-level deployment or intensity of AI tools within adopting airlines. We therefore treat enterprise AI adoption as a carrier-level technology shock and estimate its effects using variation across airline-route pairs before and after adoption, relative to comparable non-adopting carriers and routes.

For GenAI-enhanced pricing platforms, the empirical variation should similarly be interpreted as exposure to a carrier-level pricing-regime transition rather than direct observation of route-level deployment. Public vendor materials and airline statements describe these systems as being scaled gradually and experimentally, including through pilot programs and A/B testing, rather than as fully deployed across all routes at the moment of carrier adoption. This distinction is important because adoption does not necessarily imply full route-level implementation. Conditional on carrier adoption, route fixed effects, time fixed effects, and observed route characteristics, the early rollout structure makes route-level exposure less likely to be selected directly on contemporaneous route-level pricing shocks. The GenAI analysis therefore exploits variation in exposure to carrier-level vendor adoption, common-vendor overlap, and market-wide platform diffusion across route markets.

Specifically, we estimate the following specification:

$$Y_{ijt} = \alpha + \beta \cdot \mathrm{AI}_{it} + X'_{ijt}\gamma + \mu_i + \lambda_j + \delta_t + \varepsilon_{ijt}, \tag{3}$$

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. X_{ijt} includes control variables such as lagged route concentration (HHI), multimarket contact (MMC), average itinerary distance (measured by DB1B Market Distance) and its quadratic term, lagged origin and destination airport network size, the lagged number of tickets sold, and a merger indicator equal to one if airline i was involved in a merger at time t. AI_{it} is an indicator equal to one for airline-route observations associated with airline i in periods after airline i has adopted enterprise AI tools.[14] Since enterprise AI adoption occurs at the airline-time level, the coefficient β captures the average effect of enterprise AI adoption on route-level outcomes for adopting airlines, relative

[14]We use the first documented enterprise AI adoption as the treatment date because it captures the beginning of a carrier's AI-enabled operational regime. Later enterprise AI applications are typically layered onto the same data infrastructure, workflows, and decision systems. In comparison, GenAI-enhanced pricing platforms are analyzed separately because they introduce a distinct simulation and recommendation layer in pricing decisions through third-party vendors, thus representing a qualitatively different pricing paradigm.

to non-adopting airlines and pre-adoption observations.

Our baseline specifications include airline fixed effects μ_i, route fixed effects λ_j, and time (year-quarter) fixed effects δ_t to control for airline-specific factors, time-invariant route characteristics, and common temporal shocks, respectively. Identification comes from before-after changes within airline-route pairs of adopting carriers, relative to contemporaneous changes among non-adopting carriers, conditional on observed route characteristics X_{ijt} and airline, route, and time fixed effects. 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 constraint. 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 Section 6.2 and Appendix D.2.

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 industry-wide trends toward AI adoption, as opposed to short-term demand, weather, or other route-specific shocks.[15] 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 following Rosenbaum and Rubin (1983), Wooldridge (2010), and Sweeting et al. (2020). 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 instead of 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 itinerary distance, and total tickets sold. We then replace the treatment indicator in Equation 3 with the fitted adoption probability, $\widehat{\Pr(\text{AI Adoption})}_{it} = \Phi(\hat{Z}_{it}\hat{\delta})$, estimated from the first-stage probit model. The generalized regressor captures systematic differences in airlines' underlying propensity to adopt enterprise AI, and conditioning on it helps account for

[15]Note that the enterprise AI deployments we study were not accompanied by major reorganizations, bankruptcies, or 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 investment or strategy. This mitigates a potential concern that our estimated effects are driven by contemporaneous carrier-wide changes unrelated to AI adoption.

selection on observable carrier characteristics that may otherwise confound the estimated effects of AI adoption on route-level outcomes. We treat this correction as a robustness check rather than the main source of causal identification. While the two-step approach addresses selection on observables under the conditional independence assumption (CIA), it does not rule out potential bias from unobserved time-varying confounders such as concurrent managerial reorganizations or undisclosed technology initiatives. We instead rely on the staggered DID design, the parallel pre-trends evidence in Section 6.1, and the alternative fixed effects specifications presented in Section 6.2 and Appendix D.2 for causal identification.

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. The intuition behind our GenAI analysis is again analogous to retrospective airline merger studies, where the merger event occurs at the carrier level but product market impacts are estimated using route-level variation in exposure. In our setting, GenAI platform adoption is observed at the airline-time level, while route markets differ in their exposure to carrier-level adoption through common-vendor overlap and GenAI platform diffusion. 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. We interpret $GenAI_{it}$ as an indicator that carrier i has entered a GenAI-enhanced pricing regime. Because the adoption decision occurs at the carrier or vendor-platform level, while the pricing effects are measured through route-level outcomes, β_1, the coefficient on $GenAI_{it}$, represents the average route-level pricing effect associated with carriers entering the GenAI pricing regime during the early diffusion stage.

The remaining coefficients in Equation 4 capture how GenAI platform use interacts with vendor commonality and platform diffusion. Specifically, 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 can constrain identification. We argue, however, that the concentrated pattern of vendor adoption provides 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-specific fare changes or demand shocks. 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. Public documents of such GenAI rollouts emphasize gradual scaling and experimentation instead of a discrete set of publicly identifiable "GenAI routes" (see Section 7.3 for further discussion). Studying this early diffusion stage, analogous to the transition window in merger studies, allows us to estimate the average route-level pricing effect of carriers entering the GenAI pricing regime before widespread adoption potentially shifts the market toward a new pricing equilibrium.

5 Impacts of AI Adoption and Diffusion

5.1 On-Time Performance

In Table 1, columns 1–2 in Panel A 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 0.07 percentage points. This result likely reflects operational reoptimization, where airlines adopting enterprise AI prioritize minimizing extended delays and use cancellations to rebalance operations.

Panel B implements a control function specification that retains the AI adoption indicator and adds the residual from the first-stage probit model of enterprise AI adoption. Panel C instead replaces the AI adoption indicator with the fitted adoption probability from the same first-stage probit. The resulting estimates are nearly identical to those in columns 1–2 in Panel A,[16] which suggest that the results are not sensitive to how the first-stage correction is incorporated. Relative to columns 1–2 in Panel A, the two-step estimates in columns 1–2 in Panel B and Panel C are larger in magnitude, which suggests that selection on observable carrier characteristics may attenuate the estimated effects in the baseline specification.

It is worth noting that in a hub-and-spoke network, operational spillovers across connected routes imply that improvements (or disruptions) can 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

Column 3 in Panels A, B, and C of Table 1 further indicates that enterprise AI adoption can improve local air quality. The baseline estimate in column 3 shows a reduction in AQI of about 2.2 points,

[16]The first-stage results are reported in Appendix Table A3. The Wald χ^2 statistic of 313.92 indicates that observable carrier characteristics significantly predict enterprise AI adoption timing.

while the two-step specifications in columns 6 and 9 yield larger estimated reductions of about 4.0 points. Since AQI is a scaled index where lower values indicate cleaner air, the baseline estimate implies an economically significant improvement of roughly 10 percent relative to the sample mean of 22.55 shown in Appendix Table A2. 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 exposure for communities surrounding the airports.

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 A11 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.

5.3 Average Fares and Price Distribution

Our DID estimates in columns 4–6 in Table 1 indicate that enterprise AI adoption increases average airfares by 1%. The increase in average fares 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 control function estimates in columns 4–6 in Panel B and the generalized regressor estimates in columns 4–6 in Panel C closely corroborate the OLS results and preserve the same qualitative pattern across the 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 3 examines the price effects of GenAI platform use. Columns 1–3 in Panel A present the baseline estimates and indicate that carrier route observations exposed to a carrier's GenAI pricing regime exhibit average fares about 5% higher than comparable observations. 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 1, the estimates in Table 3 indicate considerably larger and more uniform price impacts.

Columns 1–3 in Panel B report the estimates from Equation 4 that incorporates the vendor commonality and platform diffusion measures. Focusing on column 1 in Panel B, the main effects of SVOI and VAII are 0.430 and -0.156, respectively, indicating that greater common-vendor overlap (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%. The effect increases with common-vendor overlap (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).

We further examine the marginal effects of *GenAI*, *SVOI*, and *VAII*, and we find them in line 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 2–3 in Panel B extend the results to the fare distribution. GenAI platform use is associated with higher fares throughout the distribution, with the marginal effects indicating increases of about 6.4% at the upper end and 4.3% at the lower end. The interaction terms show that the incremental GenAI effect varies with the vendor environment: the $GenAI \times SVOI$ coefficients are negative in both columns, with a larger negative interaction for low-end fares, while the $GenAI \times VAII$ coefficients differ across the distribution. These interaction coefficients, however, should be interpreted together with the main effects and marginal effects. The economically relevant marginal effects indicate that SVOI is positively associated with fares across the distribution, with the largest increase at the lower end. By contrast, greater platform diffusion, measured by VAII, is associated with lower fares across the distribution, again with a larger decline at the lower end. Overall, the fare-distribution results suggest that GenAI platform use raises fares, common-vendor overlap is associated with higher fares especially for low-end fares, and broader market-wide platform diffusion is associated with lower fares.

One potential concern is that the sample period for the GenAI analysis in Table 3 begins before enterprise AI adoption appears in our data, so the estimates may partly reflect comparisons with the pre-enterprise-AI period. Columns 4–6 in Table 3 address this concern by restricting the GenAI sample to 2015Q1–2025Q2, beginning with the first year of enterprise AI adoption, and we find the results to remain qualitatively similar. In particular, GenAI platform use continues to be associated with higher fares and the marginal effects of SVOI (VAII) remain positive (negative). The estimated GenAI effects appear somewhat smaller than in column 1–3 of Table 3, while the SVOI and VAII effects remain close in magnitude. These results ensure that the main GenAI findings are not driven by comparisons with the pre-enterprise-AI period.

Taken as a whole, the results from Table 3 indicate that GenAI platform use can raise prices in isolation, but its ability to foster 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, and the effect is again more pronounced at the lower end.

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, staggered AI adoption implies that time fixed effects may be estimated using both treated and untreated units. 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 staggered-adoption DID (Goodman-Bacon, 2021; Sun and Abraham, 2021; Callaway and Sant'Anna, 2021) and implement the Callaway and Sant'Anna (2021) (CS) estimator. The CS estimator compares treated observations to appropriate control groups of not-yet-treated and never-treated units, then aggregates group-time effects into overall and dynamic treatment effects. This approach is robust to treatment-effect heterogeneity and avoids using already-treated units as controls. Figure 2 plots the dynamic effects of enterprise AI adoption on airfares relative to adoption timing.[18] Pre-adoption coefficients are statistically insignificant, providing evidence consistent with parallel trends and the absence of anticipatory pricing behavior.

6.2 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 series of specifications with increasingly stringent fixed effect structures. In particular, Table 2 indicates that the estimated effect of enterprise AI adoption on fares remains 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

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

[18] The dynamic analysis focuses on enterprise AI because GenAI adoption occurs relatively late in our sample, leaving limited post-adoption periods.

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. Airline-route fixed effects absorb all time-invariant route characteristics that are specific to a given carrier on a given route. 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 2, 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 1. While 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 by persistent differences across adopting and non-adopting carriers or by gradual carrier- or route-level trends. Also note that although adoption occurs at the airline-time level, airline-by-route fixed effects do not absorb the treatment because treatment status changes over time within each carrier-route pair after adoption.

For the GenAI platform use analysis, we estimate analogous specifications with airline-by-route fixed effects, route fixed effects, and time fixed effects. This exercise is especially relevant because we do not observe route-level deployment of GenAI pricing tools. Available institutional details, including Fetcherr's A/B testing and Delta's reliance on aggregated rather than route-specific demand signals, suggest that deployment within adopting airlines is unlikely to be systematically tied to route-level pricing decisions (see Section 7.3 for further discussion). Airline-by-route fixed effects then provide an additional safeguard by absorbing time-invariant heterogeneity within each carrier-route pair. The results in Panel A of Table A9 remain consistent with those reported in Table 3, indicating that the GenAI findings are likewise robust to controlling for common time shocks and time-invariant heterogeneity at the carrier-route level.

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), there is limited variation in vendor overlap and platform diffusion. As a result, the estimates primarily capture effects among early adopters as opposed to effects once these platforms are more widely adopted. Meanwhile, the small number of vendors alleviates concerns about endogenous vendor choice, since vendor selection is essentially binary and largely driven by strategic partnerships instead of 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 of which details are discussed in D.1. The estimated signs and magnitudes in Panel B of Table A9 remain consistent with the marginal effects reported in Table 3. Overall, these results indicate that our main findings are robust to potential endogeneity in SVOI and VAII under the control function approach.

6.4 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 or GenAI platform use, and re-estimate the corresponding route-level pricing regressions. Since adoption is observed at the airline-time level, the results in Table A10 allow us to examine the impact of route-level AI adoption intensity on airline pricing. Columns 1–3 report the estimates for enterprise AI, and columns 4–6 report the corresponding estimates for GenAI platform use. Across both sets of specifications, a higher adoption share is associated with lower average fares, while the positive interaction with lagged HHI suggests that this effect becomes less negative and can turn positive on more concentrated routes. Columns 2–3 and 5–6 show similar patterns for lower-end and higher-end fares, with interaction effects that are comparable across the two ends of the fare distribution and indicate only limited effects on fare dispersion.

6.5 Interactive Factor Fixed Effects.

Another concern is that our baseline specification cannot include airline-by-time fixed effects, since these would absorb the airline-level AI adoption indicators. To allow latent common shocks to affect carriers differently over time, we re-estimate our main fare specifications using the interactive fixed effects framework of Bai (2009) on restricted subsamples (routes with more than 10,000 passengers for enterprise AI and more than 5,000 passengers for GenAI). As reported in Appendix Table A5 and discussed in Appendix D.2, enterprise AI adoption continues to raise average fares by 1.8%, with effects of 2.5% and 1.6% at the 20th and 80th percentiles, respectively. GenAI platform use is associated with a 2.1% increase in average fares and a 0.7% increase in lower-end fares, while the upper-end estimate is close to zero, indicating that our main fare results are not specific to the baseline additive fixed effects specification.

6.6 Leave-One-Airline-Out Regression.

Because most of the shared-vendor variation comes from United and JetBlue using PROS, we further re-estimate our baseline GenAI specification by excluding one airline at a time.[19] As shown in Appendix Table A6 and discussed in Appendix D.3, the estimated GenAI coefficients remain qualitatively similar to the baseline results across all specifications and continue to point to a positive relationship between GenAI platform use and average fares, with effects across the fare distribution that align with our main findings. These results suggest that, despite limited vendor overlap during this early stage of GenAI diffusion, our baseline GenAI findings do not hinge on the inclusion of any particular airline.

[19]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.

6.7 Placebo Test

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 A7 and A7b plot the respective distributions of the estimated coefficients for AI adoption and GenAI platform use, and we find both distributions approximately normal and centered around zero, indicating that our main results are unlikely to be driven by spurious treatment assignment[20].

7 Discussion

7.1 Interpreting Pricing Effects of AI Adoption

Our baseline results indicate that, while enterprise AI adoption improves service quality and local air quality, it also leads to higher average fares and lower price dispersion. 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 1 is that enterprise AI adoption improves service quality and operational reliability, which raises passengers' willingness to pay and thus airfares. To distinguish price effects from quality improvements, we estimate a nested logit demand model that includes delay and air quality as product characteristics and use the estimated market shares to calculate the markup for each observation.[21] We then regress these markup estimates on the AI adoption indicator and the same set of fixed effects. Column 1 of Table A11 reports a positive effect of enterprise AI adoption on markups, indicating that even after accounting for service quality gains, prices still rise with AI adoption. This is consistent with the price effects operating through a markup or conduct channel and supports the coordination interpretation.

Two additional tests reinforce this interpretation. First, the coefficient on AI adoption remains positive and significant after including lagged flight cancellation rates as an additional control in the baseline price regression (column 2 of Table A11), suggesting that the remaining price effect reflects increased market power rather than higher willingness to pay from improved service quality. Second, we compare residuals from the baseline price regression with those from a specification that excludes enterprise AI adoption to isolate the portion of fares attributable to AI adoption, and apply the same procedure to delay time and cancellations to construct $\Delta Delay$ and $\Delta Cancellation$. Regressing the

[20] We provide additional robustness checks that address measurement error in adoption timing and COVID-19 disruptions in E-Companion D.

[21] Details on the nested logit demand model are provided in Section 8.1.

fare residual differences on these quality measures using airline-route averages, column 3 shows that both coefficients are insignificant, indicating that the fare effects are not driven by improvements in on-time performance. Together, these results suggest that the observed price effects of enterprise AI adoption are unlikely to be driven by higher willingness to pay for improved service quality and are instead more consistent with coordination channels.

7.1.2 Capacity Reductions

Another 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 instead of changes in pricing behavior.

To assess this possibility, we directly examine whether capacity declines following enterprise AI adoption. Specifically, we regress the log number of departures and seats on AI adoption using the same baseline specification and fixed effects structure and report the results in columns 4 and 5 of Table A11. 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. Enterprise AI-based revenue management systems could raise average fares by allocating seats more efficiently across fare buckets and smoothing pricing across transactions, which may yield lower price dispersion similar to that under price coordination.

The two channels, however, differ in how price dispersion should vary with market concentration. Revenue management optimization does not generate a clear prediction along this dimension, whereas price coordination implies that incentives to undercut are weaker in more concentrated routes, where fewer competitors and repeated interaction make deviations easier to detect and punish. Appendix Figure A4 provides descriptive evidence consistent with this intuition in that more concentrated routes tend to have higher fares throughout the fare distribution. More formally, columns 1–3 of Table A10 interact the route-level share of enterprise AI adopters with lagged HHI and, as discussed in Section 6.4, show that the negative association between AI adoption share and fares weakens and may turn positive as concentration rises.[22] Because the interaction effects are similar at the lower and upper ends of the fare distribution, the estimates imply only modest changes in fare dispersion as concentration increases and provide suggestive evidence of stronger fare compression in more concentrated markets.

The two channels also differ in how low-end fares should behave across competing carriers. Under revenue management optimization, low fares reflect carrier-specific demand, booking pace, and inventory conditions, so price cuts should occur at different times across carriers. Coordination-based channels instead imply faster matching of low-fare deviations across airlines, making low fares

[22]Columns 4–6 of Table A10 report qualitatively similar estimates based on GenAI platform use intensity.

less persistent. This distinction is separate from the distributional regression results above, which show compression within an airline-route fare distribution because lower-tail fares rise more than upper-tail fares. Appendix Figure A5 provides complementary descriptive evidence at the route level. The share of low-fare observations tends to be lower when all airlines on a route have adopted enterprise AI than when none have (Figure A5a), while route-level fare dispersion, measured across competing carriers and fare observations within a route, does not fall systematically following adoption and may even increase for many routes (Figure A5b). Thus, the evidence does not support a simple revenue-management explanation in which carriers independently generate more low-fare realizations through idiosyncratic pricing adjustments. Instead, it is more consistent with enterprise AI adoption being associated with fewer low-fare realizations, even though this need not mechanically reduce overall route-level fare variability.

Overall, the evidence thus far points to higher average fares, modestly more compressed price dispersion, stronger effects on more concentrated routes, and descriptive evidence of fewer low-end fare realizations after enterprise AI adoption. These patterns are difficult to reconcile with revenue management optimization alone and are instead more consistent with tacit coordination facilitated by algorithmic pricing (Calvano et al., 2020).

7.2 Coordination Channels

To provide intuition, we develop a simple model in which GenAI platform use affects equilibrium fares through common-vendor overlap and platform diffusion. Building on the partial profit internalization framework of Bresnahan (1982) and Bernheim and Whinston (1990), we consider a differentiated Bertrand pricing game in which vendor adoption affects both the set of rivals whose profits enter an airline's objective function and the intensity of platform use in the market. The model predicts that greater same-vendor overlap raises fares, whereas broader platform diffusion lowers fares, consistent with the results in Table 3. Details are provided in E-Companion E.

As discussed in Section 2.3, the effects of same-vendor overlap and platform diffusion are theoretically ambiguous. Shared vendors may raise fares by aligning pricing responses or lower fares through stronger algorithmic competition (Harrington, 2022; Musolff, 2025). Likewise, wider platform diffusion may facilitate coordination by making rivals' behavior more predictable (Calvano et al., 2020; Wang et al., 2024b) or lower fares by improving competitive responsiveness. Our findings suggest that the alignment channel dominates for same-vendor overlap, while the competitive-responsiveness channel dominates for platform diffusion. Appendix Figure A6 summarizes this interpretation by contrasting the enterprise AI forecast-and-optimize pipeline with the GenAI simulate-and-recommend layer.

7.3 Heterogeneity by Route Size

Airlines do not disclose route- or flight-level use of GenAI-enhanced pricing tools, and public statements emphasize that these systems are integrated into broader revenue management platforms and

gradually rolled out rather than activated as discrete route-level treatments. As a result, one cannot directly observe route-level implementation or identify a set of "GenAI routes." For example, Delta 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, instead of deployment on a clearly identifiable set of routes.[23]

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 fares. In particular, higher-traffic routes involve greater passenger volume and provide a larger base over which pricing systems can affect market outcomes. This motivates two complementary heterogeneity analyses based on predetermined route size.

First, Table A12 divides routes into small and large markets using a predetermined route-size classification based on average ticket sales in 2013–2014, before any airline in our sample adopted enterprise AI. The results show that the pricing effects of both enterprise AI adoption and GenAI platform use are stronger in larger pre-AI markets. For enterprise AI, the effect on average fares is near zero in small markets but positive in large markets. Similar patterns appear across the fare distribution, with both lower- and upper-end fares increasing more in larger markets. For GenAI platform use, fare effects are positive throughout the distribution in both subsamples but substantially larger on routes with higher pre-AI ticket volume. Specifically, average fares increase by about 3.6% on small routes and 6.3% on large routes, with comparable differences at the 20th and 80th percentiles. Overall, Table A12 suggests that AI-related pricing effects are more pronounced in larger markets, where algorithmic pricing likely operates at greater scale.

Next, we estimate a continuous baseline scale specification that interacts the GenAI use indicator with $BaselineScale_j^{pre} = \log(1 + Passengers_j^{pre})$, where $Passengers_j^{pre}$ denotes total route-level passengers measured before 2023Q4, when the first GenAI-enhanced pricing platform was adopted in our sample. $BaselineScale_j^{pre}$ reflects predetermined route traffic volume and is less vulnerable to reverse causality than contemporaneous fares, OTP, or realized revenue. Table A13 reports the corresponding estimates. The interaction term between $GenAI$ and the continuous baseline scale measure is positive and precisely estimated across the fare distribution, indicating that the effect of GenAI platform use becomes more positive as predetermined route scale increases. When evaluated at the sample mean, the overall effect of $GenAI$ remains positive. The results therefore reinforce the estimates in Table A12 that, conditional on airline-level GenAI platform use, the observed fare effects are stronger on routes with greater pre-GenAI passenger volume.

[23] 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`.

8 Welfare Implications

8.1 Demand and Consumer Welfare

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. The consumer welfare calculations use the nested logit demand estimates to translate the estimated fare, delay, and AQI effects into compensating variation. The markup-deviation exercise is used separately to provide evidence on conduct. The estimation procedure is described in detail Section F in the E-Companion.

Overall, our calculations indicate substantial welfare gains from improved service quality of approximately $888.8 million, environmental benefits of roughly $424.9 million, and price-induced welfare losses of about $604.5 million, implying a net gain in consumer surplus as shown in Figure A8.[24] The additional consumer welfare loss from GenAI pricing platform adoption amounts to $199.5 million. These estimates thus suggest that AI adoption in the airline industry leads to positive net consumer surplus, but the fare effects imply a substantial consumer-to-firm transfer.

8.2 Revenue Counterfactuals

To complement the consumer welfare analysis, we compute firm-level variable revenue counterfactuals using nested logit demand estimates and Bertrand pricing first-order conditions[25]

Figure 3 presents the model-implied revenue changes attributable to AI adoption for each carrier. Because percentage changes can be sensitive to the scale of the counterfactual baseline, especially when pre-adoption revenue is relatively small, we interpret these estimates as relative model-based changes rather than accounting measures of revenue growth. For enterprise AI adoption, as shown in Panel (a), the estimated revenue effects vary substantially across carriers, with American, Delta, United, and Southwest exhibiting the largest relative change. For generative AI adoption, as shown in Panel (b), we restrict attention to the three adopting carriers, Delta, United, and JetBlue. Among the GenAI adopters, the implied revenue gains are largest for Delta and United, while JetBlue shows a smaller but still positive gain. To assess the magnitude of the simulated revenue effects, we benchmark the implied dollar changes against carriers' reported operating revenues. For Delta, for example, the estimated GenAI-related revenue gain of approximately $300 million over the post-adoption period corresponds to about 0.17 percent of Delta's cumulative operating revenue of roughly $180 billion over 2023–2025. Thus, while the model-implied relative changes in

[24] The environmental benefit estimate monetizes the estimated 2.16 AQI unit reduction over the passenger base on AI-adopted routes by combining the EPA's central VSL of $11.6 million with the concentration-response function from Pope III et al. (2002). Because AQI is a composite index whose marginal health damages depend on the pollutant driving the index, we treat this figure as approximate. Restricting the calculation to the passenger base likely understates the full environmental benefit, since improved air quality near airports may also benefit nearby residents, in which case the net consumer welfare effect of enterprise AI adoption would be even more positive.

[25] For counterfactual revenues for GenAI adoption, we use a simpler nested logit demand model without delay time and AQI. The sample includes Delta, United, American, Frontier, Southwest, and JetBlue which is consistent with Table 3. Other details on the counterfactual revenue calculations are provided in Appendix H.1.

the figure may appear large because of denominator effects, the corresponding dollar magnitudes are modest relative to carriers' total operating revenues. The revenue counterfactuals therefore suggest that AI adoption generates modest, incremental revenue improvements through pricing and revenue management channels, as opposed to large expansions in total demand.

To provide markup-based evidence on the conduct channel, we estimate a conduct index θ that captures whether enterprise AI adoption is associated with larger cost-based markup increases on airline-routes with greater pre-adoption Bertrand-implied markups. The estimate $\hat{\theta} = 0.121$ in Table A14 implies that enterprise AI adoption is associated with an additional cost-based markup increase equal to about 12.1% of the airline-route's pre-adoption Bertrand-implied markup. This can be interpreted as semi-structural evidence consistent with a conduct channel.

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, again with stronger effects 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 firms using algorithmic platforms. Policymakers have recently flagged concerns about the use of AI to personalize airfares for individual consumers, and regulators have begun examining surveillance pricing more generally, including the role of third-party pricing intermediaries that use consumer-level data and algorithmic tools.[26] Our analysis uses route-level data to examine whether AI adoption is associated with systematic changes in average fares and fare distributions, rather than individualized pricing based on consumer- or session-level characteristics. Our findings therefore speak to a different policy question in that the risks of algorithmic coordination depend not only on the presence of AI-driven pricing tools, but also on their deployment architecture. Proprietary enterprise AI, common-vendor GenAI platforms, and broader market-wide platform diffusion can have different competitive implications. In this sense, not all AI adoption is equal. Regulatory oversight should therefore focus not only on whether AI-driven pricing tools are used, but also on how they are deployed across firms, vendors, and markets. Our results also suggest that AI pricing should be evaluated in light of both its efficiency effects and its potential coordination

[26] Source: David Shepardson, "U.S. criticizes use of AI to personalize airline ticket prices, would investigate," *Reuters*, August 5, 2025. For regulatory scrutiny of surveillance pricing, see Federal Trade Commission, "FTC Surveillance Pricing Study Indicates Wide Range of Personal Data Used to Set Individualized Consumer Prices," January 17, 2025.

risks. On the firm side, relevant questions include whether pricing algorithms are reviewed before deployment, whether human oversight remains meaningful, and how AI-generated recommendations enter price-setting decisions. Consumer welfare should also be assessed more broadly, taking into account service quality, environmental externalities, fares, and coordination risks.

Our study 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 adopters 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. Lastly, access to transaction-level search and quote data could help future research distinguish changes in posted fares for a given itinerary from personalized or surveillance pricing, which is distinct from the route-level coordination channels examined in this study.

References

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). "Artificial intelligence and jobs: Evidence from online vacancies." *Journal of Labor Economics*, *40*(S1), S293–S340.

Aleksenko, S., and Miklós-Thal, J. (2026). "Third-party pricing algorithms and information sharing." Working paper.

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.

Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). "Artificial intelligence, firm growth, and product innovation." *Journal of Financial Economics*, *151*, 103745.

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." *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.

Björnerstedt, J., and Verboven, F. (2016). "Does merger simulation work? Evidence from the Swedish analgesics market." *American Economic Journal: Applied Economics*, *8*(3), 125–164.

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.

McElheran, K., Li, J. F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E. M., Foster, L. S., and Zolas, N. J. (2024). "AI adoption in America: Who, what, and where." *Journal of Economics & Management*

Strategy, *33*(2), 375–415.

Miklós-Thal, J., and Tucker, C. (2024). "AI, algorithmic pricing, and collusion." *CPI Antitrust Chronicle*, February, Vol. 2.

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.

Nevo, A. (2000). "Mergers with differentiated products: The case of the ready-to-eat cereal industry." *RAND Journal of Economics*, *31*(3), 395–421.

Pope III, C. A., Burnett, R. T., Thun, M. J., Calle, E. E., Krewski, D., Ito, K., and Thurston, G. D. (2002). "Lung cancer, cardiopulmonary mortality, and long-term exposure to fine particulate air pollution." *Journal of the American Medical Association*, *287*(9), 1132–1141.

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.

Prince, J. T., and Simon, D. H. (2017). "The impact of mergers on quality provision: Evidence from the airline industry." *Journal of Industrial Economics*, *65*(2), 336–362.

Rosenbaum, P. R., and Rubin, D. B. (1983). "The central role of the propensity score in observational studies for causal effects." *Biometrika*, *70*(1), 41–55.

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.

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.

Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data*. MIT Press, 2nd edn.

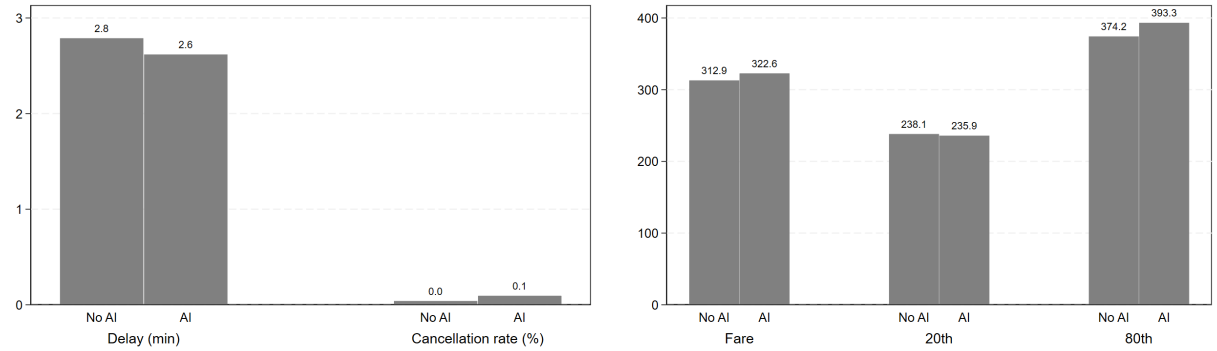
Zhang, S., Mehta, N., Singh, P. V., and Srinivasan, K. (2021). "Frontiers: Can an artificial intelligence algorithm mitigate racial economic inequality? An analysis in the context of Airbnb." *Marketing Science*, *40*(5), 813–820.

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

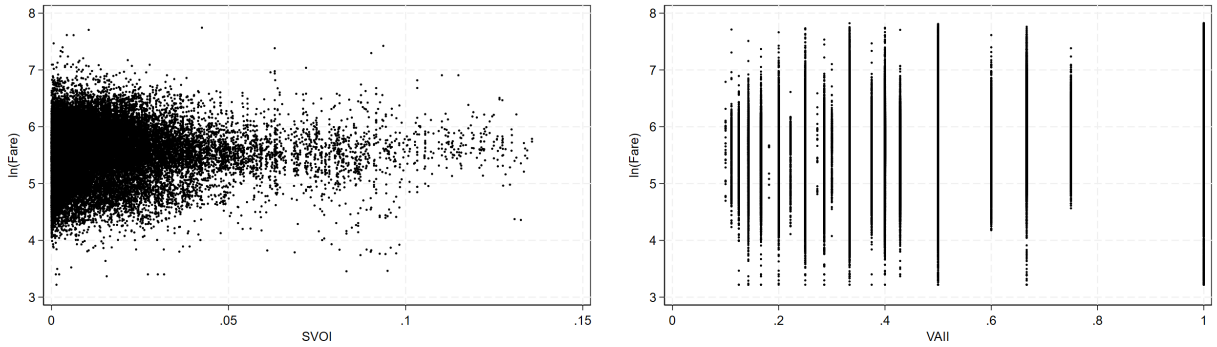
(a) OTP by Enterprise AI Adoption

(b) Fare by Enterprise AI Adoption



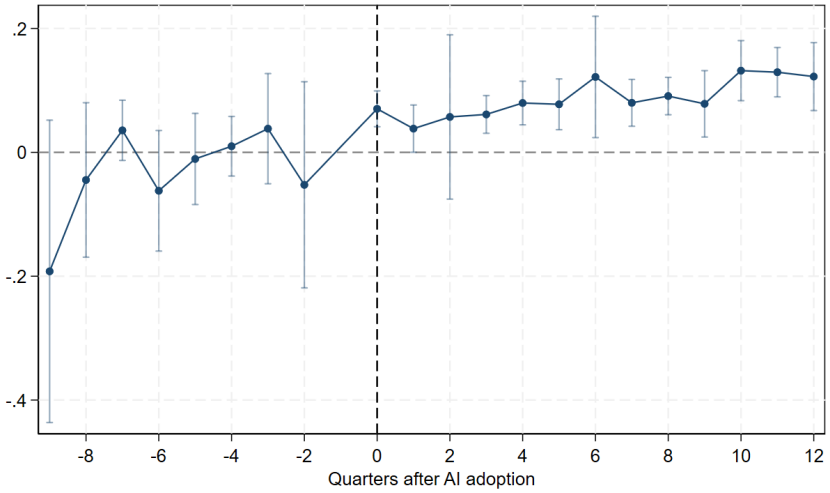
(c) SVOI vs. Fares

(d) VAII vs. Fares



Note: Sample periods are 2013Q1–2023Q3 for Panels (a) and (b) and 2013Q1–2025Q2 for Panels (c) and (d). Panels (a) and (b) display unconditional means by enterprise AI adoption status, where *AI* denotes airline-route-quarter observations in which the carrier had adopted enterprise AI by time t, and *No AI* denotes observations prior to adoption or observations for non-adopting carriers. Delay time and cancellation rates exclude weather-related disruptions. Fares are in nominal dollars. The 20th and 80th percentile fares are computed within each airline–route–quarter cell and serve as proxies for leisure and business traveler segments, respectively. Panels (c) and (d) are scatter plots of log average fares against SVOI and VAII, respectively, and include only observations with positive index values. SVOI measures the share of weighted airline pairs on a route that rely on the same GenAI pricing vendor, and VAII measures the share of airlines on a route using any third-party GenAI pricing vendor.

Figure 2: Robustness Check: Dynamic Treatment Effects of Enterprise AI Adoption on Airfares under Staggered Adoption

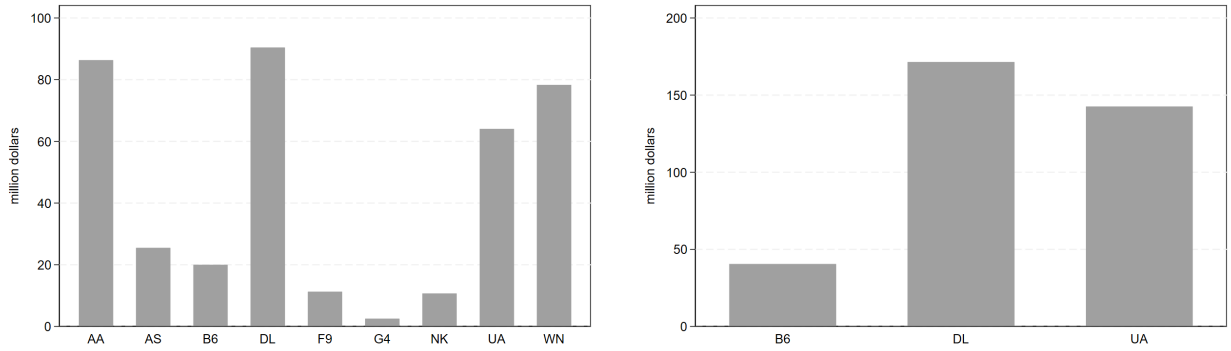


Note: This model restricts the sample to periods through 2021 and thus excludes the post-2021 period when all airlines had adopted enterprise AI. In the first stage, lagged HHI and MMC are regressed on the lagged number of LCCs, the sum of competitors' prices, passengers, and distances in the market. The corresponding fitted values are then used in the second-stage regression.

Figure 3: Counterfactual Variable Revenue Effects of AI Adoption

(a) Enterprise AI Adoption

(b) GenAI Adoption



Note: The sample period is 2013Q1–2023Q3 for Panel (a) and 2013Q1–2025Q2 for Panel (b). Revenue is measured as the average quarterly change in revenue after AI adoption; Panel (a) corresponds to enterprise AI adoption and Panel (b) corresponds to GenAI adoption. For Panel (b), the sample is restricted to the three GenAI adopters, Delta, United, and JetBlue, so the estimates capture direct adoption effects rather than market-level spillovers for non-adopting carriers. All dollar values are reported in millions of dollars. Variable revenue counterfactuals are computed under Bertrand-Nash conduct, following the standard merger simulation approach of Nevo (2000) and Björnerstedt and Verboven (2016). The Bertrand-Nash assumption is applied symmetrically to the actual and counterfactual scenarios, so the estimates capture the demand-side contribution of AI adoption through changes in quantities and prices, holding conduct fixed.

Table 1: Enterprise AI Adoption, Operational Performance, Air Quality, and Airfare Distribution

	(1) Delay Time	(2) Cancellation	(3) AQI	(4) lnFare	(5) lnFare20	(6) lnFare80
Panel A. Baseline DID						
AI Adoption	-0.394***	0.070***	-2.161***	0.010***	0.025***	0.008***
	(0.056)	(0.004)	(0.048)	(0.001)	(0.001)	(0.001)
Observations	3,304,316	3,304,316	2,682,106	3,304,316	3,304,316	3,304,316
Panel B. Control Function						
AI Adoption	-0.905***	0.171***	-3.996***	0.009***	0.033***	0.009***
	(0.096)	(0.007)	(0.084)	(0.001)	(0.002)	(0.002)
Observations	3,304,316	3,304,316	2,682,106	3,302,657	3,302,657	3,302,657
Panel C. Generalized Regressor						
Pr(AI $\widehat{\text{Adoption}}$)	-0.908***	0.171***	-3.989***	0.010***	0.034***	0.010***
	(0.096)	(0.007)	(0.083)	(0.001)	(0.002)	(0.002)
Observations	3,304,316	3,304,316	2,682,106	3,302,657	3,302,657	3,302,657

Note: Sample period is 2013Q1–2023Q3. Columns 1–3 report effects on operational performance and air quality; columns 4–6 report effects on the airfare distribution. *Cancellation* is measured in percentage points. *AQI* is the average AQI weighted by fuel expenses multiplied by passenger shares. Panel B implements a control function approach by including the AI adoption indicator together with the residual from the first-stage probit model of enterprise AI adoption; the residual coefficient is omitted for brevity. Panel C implements the generalized regressor approach by replacing the AI adoption indicator with $Pr(AI\ \widehat{Adoption})$, the fitted probability estimated from the same first-stage probit. First-stage results are reported in Appendix Table A3. All specifications include the same set of controls, airline fixed effects, route fixed effects, and time fixed effects. Control variables include lagged HHI, MMC, average itinerary distance and its quadratic term, lagged origin and destination network size, the lagged number of tickets sold, weather, fuel price, and a merger indicator. Instrumental variables for lagged 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 2: Robustness Check: Impact of Enterprise AI Adoption on Airfare Distribution (Alternative Fixed Effects)

	Route-Specific Trends			Airline-Specific Trends			Airline-Route FE		
	(1) lnFare	(2) lnFare20	(3) lnFare80	(4) lnFare	(5) lnFare20	(6) lnFare80	(7) lnFare	(8) lnFare20	(9) lnFare80
AI Adoption	0.011***	0.026***	0.009***	0.023***	0.036***	0.019***	0.018***	0.022***	0.018***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	3,304,316	3,304,316	3,304,316	3,304,316	3,304,316	3,304,316	3,296,395	3,296,395	3,296,395
Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y
Airline FE	Y	Y	Y	Y	Y	Y	N	N	N
Route FE	Y	Y	Y	Y	Y	Y	N	N	N
Airline-Route FE	N	N	N	N	N	N	Y	Y	Y
Route FE × Trend	Y	Y	Y	N	N	N	N	N	N
Airline FE × Trend	N	N	N	Y	Y	Y	N	N	N
Time FE	Y	Y	Y	Y	Y	Y	Y	Y	Y

Note: Sample period is 2013Q1–2023Q3. All specifications include the same set of controls as in Table 1. Standard errors in parentheses are clustered by route. *** p<0.01, ** p<0.05, * p<0.1.

Table 3: Impact of GenAI Platform Use and Vendor Sharing on Airfare Distribution

	Full Sample			Post-Enterprise-AI Period		
	(1) lnFare	(2) lnFare20	(3) lnFare80	(4) lnFare	(5) lnFare20	(6) lnFare80
Panel A. Baseline GenAI Effects						
GenAI	0.050***	0.037***	0.041***	0.038***	0.027***	0.028***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	3,687,969	3,687,969	3,687,969	3,207,581	3,207,581	3,207,581
Panel B. Vendor Sharing and Platform Diffusion						
GenAI	0.077***	0.039***	0.063***	0.054***	0.019***	0.040***
	(0.002)	(0.003)	(0.003)	(0.002)	(0.003)	(0.003)
SVOI	0.430**	1.945***	0.744***	0.553***	2.211***	0.839***
	(0.176)	(0.221)	(0.202)	(0.172)	(0.219)	(0.197)
VAII	-0.156***	-0.275***	-0.134***	-0.169***	-0.292***	-0.145***
	(0.004)	(0.005)	(0.005)	(0.004)	(0.005)	(0.005)
GenAI × SVOI	0.209	-1.714***	-0.660***	-0.052	-2.032***	-0.917***
	(0.185)	(0.222)	(0.193)	(0.186)	(0.228)	(0.194)
GenAI × VAII	-0.025***	0.044***	-0.020***	-0.002	0.065***	0.002
	(0.005)	(0.006)	(0.005)	(0.005)	(0.006)	(0.005)
Marginal Effects						
GenAI	0.075***	0.042***	0.062***	0.053***	0.025***	0.040***
	(0.002)	(0.003)	(0.002)	(0.002)	(0.003)	(0.002)
SVOI	0.446***	1.812***	0.693***	0.548***	2.030***	0.758***
	(0.164)	(0.208)	(0.190)	(0.158)	(0.204)	(0.183)
VAII	-0.158***	-0.272***	-0.136***	-0.169***	-0.287***	-0.145***
	(0.004)	(0.005)	(0.005)	(0.004)	(0.005)	(0.005)
Observations	3,687,969	3,687,969	3,687,969	3,207,581	3,207,581	3,207,581
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: The full sample covers 2013Q1–2025Q2, while the post-enterprise-AI period covers 2015Q1–2025Q2. The sample includes Delta, United, American, Frontier, Southwest, and JetBlue. All specifications include the same set of controls as in Table 1. Standard errors in parentheses are clustered by route. *** p<0.01, ** p<0.05, * p<0.1.

E-Companion

Artificial Intelligence, Operational Efficiency, and Price Coordination

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 industry and corporate sources.[27]

American Airlines

- **2016**: Introduced AI in baggage-handling operations to improve efficiency and reduce mishandled luggage incidents.[28]
- **2020**: Expanded AI use in flight operations and crew scheduling to improve the efficiency of disruption management.[29]
- **2021**: Adopted "Smart Gating" technology which uses AI to route aircraft to the nearest available gate.[30].
- **2022**: Partnered with Microsoft to utilize Azure as the preferred cloud platform to enhance operations with AI and machine learning.[31]

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.[32]
- **2019**: Launched the *ConnectionSaver* program, an AI-based tool that predicts whether connecting passengers can make their flights and adjusts gate holds accordingly.[33]

[27] We also use Form 41 financial data and Lightcast labor data to cross validate adoption timing and construct carrier level controls, and BTS T-100 Segment data to measure departures and seat capacity in the alternative-channel analysis.

[28] *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/.

[29] *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.

[30] *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.

[31] *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-customers-and

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

[33] *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.

- **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.[34]

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.[35]
- **2023**: Began experiments with AI-driven pricing optimization in revenue-management systems.[36]

Alaska Airlines

- **2019**: Began using the Flyways AI platform to optimize air-traffic flow and improve flight efficiency.[37]

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.

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=Fetcherr (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).

[34] *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/.

[35] *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/.

[36] *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.

[37] *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-traffic-flow-wit

The (unweighted) SVOI for this route is given by

$$\text{SVOI}^{\mathbf{A}}_{ijt} = 1/3, \quad \text{SVOI}^{\mathbf{B}}_{ijt} = 1, \quad \text{SVOI}^{\mathbf{C}}_{ijt} = 0, \quad \text{SVOI}^{\mathbf{D}}_{ijt} = 0.^{38}$$

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 $\Omega_{ii'} \in [0, 1]$.
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 Descriptive Patterns

Columns 1–3 of Table A2 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 for routes served by AI-adopting carriers than routes served by non-adopting carriers, and the gap appears more

[38] 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.

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 difference 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 largely 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 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. Meanwhile, 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 in the main text 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 shows no strong association between the diffusion of GenAI-enhanced pricing platforms and fare levels.[39]

D Additional Robustness Checks

D.1 Control function

In control function approach, we estimate the following first-stage regressions:

$$\text{SVOI}_{ijt} = \beta_0 + \beta_1 \frac{1}{\left|\mathcal{J}_{o(j),t}\right|} \sum_{(i',j') \in \mathcal{J}_{o(j),t}} \text{SVOI}_{i'j',t-1} + \beta_2 \frac{1}{\left|\mathcal{J}_{d(j),t}\right|} \sum_{(i',j') \in \mathcal{J}_{d(j),t}} \text{SVOI}_{i'j',t-1} + \mathbb{X}'_{ijt}\theta$$

$$\text{VAII}_{jt} = \alpha_0 + \alpha_1 \sum_{i' \neq i, i \in \mathcal{J}_t} GenAI_{it} + \mathbb{X}'_{ijt}\gamma,$$

where $\mathcal{J}_{o(j),t}$ ($\mathcal{J}_{d(j),t}$) is the set of airline-route observations sharing the same origin (destination) airport as route j at time t. The set of control variables $\mathbb{X}$ includes lagged administrative expenses, reservation and sales expenses, and line service expenses weighted by the route-level ticket share for each airline.

In the second stage, we exclude the interaction terms and re-estimate Equation 4, while retaining *SVOI* and *VAII* and including the residuals from the *SVOI* and *VAII* first-stage regressions as additional controls.

D.2 Interactive Factor Fixed Effects

One concern in our empirical setting is that the baseline specification cannot include airline-by-time fixed effects, since these would absorb airline-level AI adoption indicators. As a result, the

[39] Additionally, Figure A2 visualizes the evolution of airfares, market concentration, OTP, and AQI, while Figure A3 plots the distribution of SVOI and VAII.

estimates may still reflect airline-specific shocks that evolve over time and coincide with adoption, such as managerial, organizational, or technological changes. For GenAI platform use, the relevant adoption decision occurs at the carrier or vendor-platform level, while pricing effects may vary across route markets as airlines incorporate these tools into existing revenue management systems. The estimates may therefore also reflect latent time-varying factors that shape both platform adoption and route-level pricing outcomes.

To mitigate these concerns, we re-estimate our main fare specifications using the interactive fixed effects framework of Bai (2009), which augments the standard additive fixed effects model with a latent factor structure that allows unobserved common shocks to affect observational units differently over time. This approach is well suited to our setting because unobserved shocks may coincide with enterprise AI or GenAI adoption and also affect pricing, yet remain unabsorbed by airline, route, and time fixed effects if those shocks affect carriers differently over time. The interactive fixed effects model helps address this concern by allowing latent common factors to have heterogeneous effects across carriers over time. Because the estimator is computationally intensive, we implement it on restricted subsamples, limiting the enterprise AI sample to routes with more than 10,000 passengers and the GenAI sample to routes with more than 5,000 passengers.

Table A5 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 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 point to higher average fares following GenAI platform adoption, although the distributional effects are more mixed than in Table 3. Overall, the interactive fixed effects estimates in Table A5 confirm that our main price results are not specific to the baseline additive fixed effects specification.

D.3 Leave-One-Airline-Out Regression

Table A6 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.[40] This design allows us to assess whether the estimated relationship between GenAI platform use and airfare distribution is driven disproportionately by any single carrier, which is particularly relevant in our setting as most of the shared-vendor variation comes from United and JetBlue using PROS. 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 effects across the fare distribution that are generally in line with our main findings in Table 3. These results therefore suggest that, despite the limited vendor overlap in this early stage of GenAI diffusion, our baseline GenAI findings do not hinge on the inclusion of any particular airline in the sample.

[40]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.

D.4 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 A7, 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.

D.5 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 A8, remain robust.

E Theoretical Framework for GenAI Pricing and Vendor Adoption

In this appendix 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 to capture the incentives created by common ownership.

E.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 outline the model assumptions before discussing the implications for vendor adoption and sharing. Assumption 1 ensures a well-behaved best-response mapping, so each profit maximization problem has a solution.

Assumption 1 (Best response function) For each firm i, the demand function $q_i(\mathbf{p}; m)$ is twice continuously differentiable on $\mathcal{P}^N$ for each $m \in [0, 1]$. The effective profit function $\tilde{\pi}_i(\mathbf{p}, m)$ 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 strategic 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 $i \neq i'$,

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

We also introduce an assumption that captures the idea that wider 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.

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

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

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

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 (Milgrom and Shannon 1994; Topkis 1998), 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 $\mathbf{p}^{\max}(m, v, \theta)$.

E.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 Milgrom and Shannon (1994) and Topkis (1998), 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.

F Demand Model

To quantify the net welfare impact on consumers, we estimate a nested logit demand model for airline tickets that includes airfares, delay time, and AQI. In particular, consumers first choose

between two nests, major airlines and LCCs, and then choose an airline within the selected nest. We specify the indirect utility of passenger k choosing airline i on route j at time t as:

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

where $\mathbb{X}_{ijt}$ denotes airline and route characteristics, ς_{ig} is consumer i's preference toward nest g, and ϵ_{ijkt} captures unobserved consumer heterogeneity. Assuming that ϵ_{ijkt} has a Type I extreme value distribution, we follow Berry (1994) and estimate the corresponding demand function as:

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

where the market size is defined as the geometric mean of origin and destination county populations divided by 20, which scales the population base to a plausible quarterly air-travel market size, $\ln(s_{j|g})$ is within-nest market share, 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, within-nest market share, delay time, and AQI.[41] The instruments include (1) the sum of competitors' characteristics on the route and its quadratic term, (2) the airline's average route distance in other markets and its quadratic term, and (3) the average precipitation competitors experience on other competing routes.

Following the market share inversion method of Berry (1994), we compute the consumer 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 Table 1 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, $\hat{v}^{cf}_{ijt}$, 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 consumer surplus gains from improved on-time performance and air quality follow the same approach, with counterfactual utilities computed by adding 0.39 minutes to delay time and 2.16 units to AQI for post-adoption observations. Similarly, we calculate the counterfactual fare without GenAI adoption as $\hat{v}_{ijt}\,(p_{ijt}/1.0513)$ using the estimate in Table 3.

G Model Identification

We use the demand estimates from Section 8.1 primarily for welfare decomposition, not to recover the causal effect of AI adoption, which is instead identified through the reduced-form DID design in Section 4 and supported by event study pre-trend tests, staggered-adoption estimators, interactive-factor models, and placebo exercises. The nested logit model provides the preference parameters necessary to convert the reduced-form estimates, including a 1% fare increase, a 0.39-minute delay reduction, and a 2.16-unit AQI improvement, into dollar-valued compensating variation. Because AI adoption affects fares, delays, and air quality simultaneously, a fare-only analysis would overstate consumer harm, while an operations-only analysis would ignore the fare effects. The estimates of α, β_1, and β_2 thus allow us to produce a welfare estimate that reflects consumers' combined exposure to changes in fares, service quality, and environmental outcomes.

[41] E-Companion G provides the relevant identification discussion for the nested logit model, including potential endogeneity concerns, excluded instruments, and identifying assumptions.

Following Berry (1994), consumers first decide whether to fly and, conditional on flying, choose among airlines serving a non-directional airport-pair market. We define nests by carrier type, separating major carriers from low-cost carriers. The estimation equation is given by

$$\ln\left(\frac{s_{ijt}}{s_{0jt}}\right) = \alpha\, p_{ijt} + \beta_1\, \text{Delay}_{ijt} + \beta_2\, \text{AQI}_{ijt} + X'_{ijt}\theta + \sigma \ln(s_{ij|g,t}) + \mu_i + \lambda_j + \delta_t + \xi_{ijt}, \qquad \text{(A1)}$$

where s_{ijt} is airline i's market share on route j at time t, s_{0jt} is the outside-option share defined using the geometric mean of origin and destination county populations divided by 20, which scales the population base to a plausible quarterly air-travel market size, p_{ijt} is average fare, Delay_{ijt} is non-weather delay time, AQI_{ijt} is the airline-route-time air quality exposure measure described below, and $\ln(s_{ij|g,t})$ is airline i's within-nest share among carriers of the same type. Airline fixed effects μ_i, route fixed effects λ_j, and time fixed effects δ_t absorb persistent carrier heterogeneity, time-invariant route characteristics, and common aggregate shocks, respectively.

The error term ξ_{ijt} captures unobserved demand factors, including business travel intensity, schedule convenience, loyalty program effects, carrier reputation, advertising or fare sale activity, and local economic activity. The identification challenge is that fares, delay time, AQI, and within-nest shares can all move with ξ_{ijt}, either because airlines respond to demand when setting prices and operational priorities, or because the same route and airport characteristics that affect demand also affect congestion and air pollution. We use instrumental variables for each endogenous regressor and discuss the identifying assumptions in turn.

G.1 Price

Fares are set in equilibrium. Airlines observe demand conditions when setting prices, so routes with stronger unobserved demand will tend to have both higher fares and higher market shares. Treating price as exogenous in Equation A1 would bias $\hat{\alpha}$ toward zero, making consumers appear less price-sensitive and inflating the estimated dollar cost of the AI-driven fare increase. An attenuated $\hat{\alpha}$ would thus overstate the compensating variation associated with a given fare change.

We use two sets of excluded instruments. The first is the sum of competitors' product characteristics on the same route and its quadratic term in the spirit of BLP. Following Berry (1994), rival product characteristics shift the competitive environment and therefore equilibrium markups without directly entering the focal airline's own demand. A route with more competitors or stronger rival products faces more downward pressure on markups, while a route with weaker rival service supports higher markups. Conditional on airline, route, and time fixed effects, competitor characteristics should be orthogonal to ξ_{ijt} for the focal airline. The main threat to this exclusion restriction is that competitor product choices may respond to the same demand conditions as the focal airline's choices. For example, rivals may expand service on a route because unobserved demand is strong. The fixed effects absorb persistent cross-carrier correlations in network strategy, and the instruments are constructed from rival attributes rather than the focal airline's realized demand. This alleviates, but does not fully eliminate, the concern.

The second instrument is the airline's average route distance in other markets and its quadratic term. Network-average distance captures cost and operational scale conditions, including fuel burn, aircraft deployment, and crew scheduling complexity, that feed into the carrier's overall cost structure and system-wide pricing without being driven by demand on any particular focal route. The identifying assumption is that, after conditioning on airline fixed effects and route characteristics, the carrier's distance profile elsewhere does not directly shift passengers' unobserved preference for that airline on the focal route. A potential violation would arise if broad network expansion into

longer haul markets were accompanied by brand or service improvements that also raise demand on shorter routes. Airline fixed effects absorb much of the persistent correlation between network composition and carrier quality, making this concern less acute, though the exclusion restriction remains an assumption rather than a testable restriction.

G.2 Delay Time

The coefficient β_1 captures passengers' disutility from delays. Non-weather delay is potentially endogenous for two distinct reasons. First, high-demand routes may generate more passenger volume and more complex airport operations, which can increase delays independent of the carrier's own operational effort. A positive ξ_{ijt} shock that attracts more passengers may therefore also worsen on-time performance and induce a positive correlation between delay and the error term that biases $\hat{\beta}_1$ toward zero. This would attenuate the estimated delay disutility, and the welfare gain from AI-driven delay reductions would be understated. Second, airlines may concentrate operational investment on high demand routes, which can further obscure the relationship between delays and demand.

We instrument using the average precipitation that competitors experience on their *other* routes. The first-stage relationship operates through airline network structure. Adverse weather on a competing carrier's other routes can generate cascading disruptions, including aircraft repositioning problems, crew misconnections, and increased taxi or holding times at shared hubs, that propagate through the broader network and affect operational performance on related routes. Note that we are not using weather on the *focal* route, which would directly affect both demand and operations. Rather, competitor weather elsewhere creates exogenous variation in network-wide congestion and capacity that shifts the focal airline's operational environment without independently shifting passengers' underlying preference for the focal airline-route product. For example, heavy precipitation on United's other routes may disrupt United's aircraft and crew schedules in ways that spill over to shared airports, but it should not directly affect passengers' unobserved preference for Delta on an unrelated route.

G.3 Local Air Quality

The coefficient β_2 measures passengers' valuation of local air quality. The variable AQI_{ijt} is constructed from EPA Air Quality System monitor readings at the nearest monitoring station to each airport, weighted by fuel expenses multiplied by passenger shares at the airline-route-time level. The weighting is appropriate for welfare analysis because passengers' effective pollution exposure depends on where airlines fly and the scale of their operations at each airport pair. However, this weighting also introduces endogeneity. Fuel expenditure and passenger volume are outcomes of airline operations that respond to demand. On high-demand routes, airlines operate more flights, burn more fuel, and contribute more to local emissions and congestion, which means AQI_{ijt} is likely positively correlated with ξ_{ijt} through the weighting mechanism itself, even though the underlying EPA readings come from externally sited monitors unrelated to individual carriers' decisions.

There is also a second, separate source of concern. Airport-level pollution reflects urban density, industrial composition, road traffic, and local economic activity, all of which may also affect air travel demand. Routes connecting large metropolitan areas tend to face stronger underlying demand and worse baseline air quality simultaneously. Route fixed effects absorb the time-invariant component of this correlation, but time-varying local economic shocks could still induce correlation between AQI_{ijt} and ξ_{ijt} within a route over time.

To address both concerns, we again use competitor precipitation on other routes as an instrument. Beyond its relevance for delay, weather-driven changes in competitors' flight activity alter aircraft movements, ground operations, and auxiliary power use at connected airports, which shifts locally generated emissions and thus measured AQI at the origin and destination airports of the focal route. The exclusion restriction is that competitor weather shocks affect the focal route through operational disruptions and flight activity, rather than through passengers' demand for the focal airline-route product. This is the same identifying assumption used for the delay instrume

G.4 Within-Nest Share

The nesting parameter σ governs the degree of substitution within carrier-type groups relative to substitution across groups. A larger σ means passengers are more likely to switch to another major carrier than to an LCC when their preferred option becomes less attractive, and vice versa. The parameter is identified from how within-nest shares vary relative to overall shares across airlines, routes, and time. The within-nest share $\ln(s_{ij|g,t})$ is endogenous by construction. It is a transformation of the same market shares on the left-hand side of Equation A1, so a positive demand shock to airline i on route j simultaneously raises both sides. Ignoring this would bias $\hat{\sigma}$ upward: demand shocks that raise a carrier's share also raise its within-nest share, making within-group substitution appear stronger than it is.

We instrument for $\ln(s_{ij|g,t})$ using the same excluded instruments as for price. Competitor characteristics and other-market distance shift rival prices and product characteristics within a nest, through rival cost differences and the relative attractiveness of competing airline products, without entering the focal airline's own demand shock. These instruments therefore shift within-nest shares through variation in the passenger choice set, which helps identify σ separately from unobserved demand.

G.5 Welfare Decomposition

To compute compensating variation, we compare predicted mean utility under observed outcomes with counterfactual utility in which the AI-induced changes are removed. For enterprise AI, this means reducing post-adoption average fares by 1.01%, increasing delay time by 0.39 minutes, and increasing AQI by 2.16 units, based on the DID estimates in Table 1. For each channel, the utility difference is divided by $\hat{\alpha}$ to express it in dollars. For GenAI platform use, we reduce average fares by 5.13% using the estimate in Table 3 and apply the same conversion. Because GenAI platform adoption occurs at the carrier level and diffuses gradually across route markets, the GenAI welfare figure captures the consumer surplus effect of the observed carrier-level fare changes during the early adoption period.

The welfare calculation rests on two distinct layers of identification. The first is the reduced-form identification of the AI treatment effects, which we discuss at length in the main text. The second is the structural identification of α, β_1, β_2, and σ through the instrumental variables strategy described above. These two layers are independent: the DID design identifies how AI changes fares, delay, and AQI, while the demand model identifies how passengers value those changes. Combining them yields a welfare decomposition that cannot be obtained from reduced-form estimates alone.

G.6 Remaining Concerns

Several caveats remain. First, weak instruments would make $\hat{\alpha}$, $\hat{\beta}_1$, $\hat{\beta}_2$, and $\hat{\sigma}$ unreliable and propagate directly into the welfare estimates. We report first-stage F-statistics for each endogenous variable. Second, despite using rival and other-market variation throughout, correlated network

shocks or broad carrier strategy changes could still affect the instruments. The fixed effect structure absorbs most persistent confounders, but time-varying carrier- or route-specific shocks that correlate with both the instruments and ξijt remain a residual concern. Third, defining nests by major versus LCC carrier type captures an important dimension of product differentiation in U.S. airline markets, but passengers may also differentiate along nonstop versus connecting service, departure time, airport convenience, and frequent flyer status. The welfare estimates are therefore conditional on the maintained nesting assumption and could differ under alternative nest definitions. Finally, the welfare counterfactual captures partial equilibrium outcomes by evaluating consumer surplus after removing the AI-induced changes in fares, delay, and AQI, while holding the demand system and market structure fixed. It does not re-solve airlines' equilibrium pricing, capacity, or adoption decisions. The estimates are therefore useful for quantifying the surplus implications of the observed treatment effects, but they should not be interpreted as predictions of long-run outcomes under a different technology regime.

H Supply Side of the Structural Model

H.1 Revenue Counterfactuals

We provide details on computing counterfactual revenues for each airline without enterprise AI and GenAI using the supply side of the structural model used for welfare analysis. For each airline i on route j at time t, we recover marginal cost as

$$\hat{c}_{ijt} = p_{ijt} - \mu^B_{ijt},$$

where the Bertrand markup μ^B_{ijt} is given by

$$\mu^B_{ijt} = -\frac{s_{ijt}}{\partial s_{ijt}/\partial p_{ijt}},$$

and the own-price derivative under nested logit is

$$\frac{\partial s_{ijt}}{\partial p_{ijt}} = \hat{\alpha}\, s_{ijt}\left[\frac{\hat{\sigma}}{1-\hat{\sigma}}\, s_{ij|g,t}\big(1 - s_{ij|g,t}\big) + (1 - s_{ijt})\right],$$

where $\hat{\alpha}$ is the estimated price coefficient, $\hat{\sigma}$ is the nesting parameter, s_{ijt} is the overall market share, and $s_{ij|g,t}$ is the within-nest market share, all taken from the demand estimates in Section 8.1. Because the demand system treats each carrier-route-quarter as a single differentiated product (its average fare), the ownership matrix is diagonal and the Bertrand markup is given exactly by the scalar own-price expression above. The markup estimates are therefore conditional on this product definition and should be interpreted as applying to average markup levels rather than to within-carrier fare-class dispersion. Variable profit for airline i on route j at time t is then

$$\pi_{ijt} = (p_{ijt} - \hat{c}_{ijt})\, q_{ijt} = \mu^B_{ijt}\, q_{ijt},$$

and revenue is analogously defined as

$$R_{ijt} = p_{ijt}\, q_{ijt}.$$

Counterfactual revenues in the absence of AI adoption are computed by replacing observed prices and market shares with their simulated no-AI counterparts, following the same counterfactual price

and quantity construction used for the consumer welfare calculations.

H.2 Markup Deviation Test

This appendix provides additional details on the construction of the conduct measure and the estimation procedure.

First, we construct a smoothed marginal-cost measure by projecting the Bertrand-implied marginal cost, obtained from the estimated demand system, onto observed cost shifters and fixed effects. Let

$$\tilde{c}_{ijt} = p_{ijt} - \mu^B_{ijt},$$

where μ^B_{ijt} is the Bertrand-implied markup obtained from the estimated demand system. We project $\tilde{c}_{ijt}$ onto observed cost shifters and fixed effects:

$$\tilde{c}_{ijt} = \delta_0 + \delta_1 FuelCost_{ijt} + \delta_2 LandingFee_{jt} + \eta_i + \nu_j + \xi_t + u_{ijt}.$$

The fitted component of this regression defines the estimated marginal cost:

$$\hat{c}_{ijt} = \widehat{\delta}_0 + \widehat{\delta}_1 FuelCost_{ijt} + \widehat{\delta}_2 LandingFee_{jt} + \widehat{\eta}_i + \widehat{\nu}_j + \widehat{\xi}_t.$$

Using fitted values rather than $\tilde{c}_{ijt}$ directly smooths out residual markup variation in the Bertrand-implied measure. The resulting cost-based markup remains conditional on the estimated demand system and the Bertrand markup calculation, and is defined as:

$$m_{ijt} = p_{ijt} - \hat{c}_{ijt}.$$

Second, we construct the pre-adoption average Bertrand markup for each airline-route pair as

$$\bar{\mu}^B_{ij,pre} = \frac{1}{|T_{ij,pre}|} \sum_{t \in T_{ij,pre}} \mu^B_{ijt},$$

where $T_{ij,pre}$ denotes the set of quarters before airline i adopted enterprise AI on route j. We use the pre-adoption average markup because estimates in Table A11 suggest that μ^B_{ijt} is itself affected by AI adoption, which would introduce a bad-control bias.

Third, we estimate the following conduct regression:

$$m_{ijt} = \alpha_0 + \theta \cdot AI_{it} \cdot \bar{\mu}^B_{ij,pre} + \beta\, \bar{\mu}^B_{ij,pre} + \gamma' X_{ijt} + \psi_{it} + \lambda_j + \varepsilon_{ijt}, \tag{A2}$$

where m_{ijt} is the cost-based markup, AI_{it} is the enterprise AI adoption indicator, ψ_{it} denotes airline-by-time fixed effects, λ_j denotes route fixed effects, and X_{ijt} contains the same controls as in the baseline pricing regressions. The airline-by-time fixed effects absorb the level effect of AI_{it}, so the conduct channel is identified from the interaction term: θ measures whether, within a given airline-quarter, AI adoption is associated with larger cost-based markup increases on routes with greater pre-adoption Bertrand-implied markups. Because $\bar{\mu}^B_{ij,pre}$ varies across routes within an airline-quarter, both its main effect and the interaction remain identified under this specification.

In Table A14, we find that the estimated coefficient $\hat{\theta}$ is positive, implying that enterprise AI adoption raises markups above the Bertrand benchmark. The estimate $\hat{\theta} = 0.121$ implies that enterprise AI adoption is associated with an additional cost-based markup increase equal to about 12.1% of the airline-route's pre-adoption Bertrand-implied markup. We interpret this as semi-structural evidence consistent with a conduct channel, rather than as a structural estimate

of movement along a $[0, 1]$ Bertrand-to-collusion scale. The magnitude is broadly in line with the proportional margin increases documented in other algorithmic-pricing settings; for example, Assad et al. (2024) report margin increases of roughly 15% in the German retail gasoline market.

One point to note is that markup-deviation estimate exceeds the reduced-form effect in Table A11, column 1, because the conduct channel does not load fully onto the demand-system-fitted Bertrand markup. The reduced-form regression uses the Bertrand-implied markup, which is mechanically constrained by the demand-system fit; the conduct regression uses the cost-based markup $m_{ijt} = p_{ijt} - \hat{c}_{ijt}$, allowing θ to capture deviations from the Bertrand benchmark.

Figure A1: AI Adoption Timelines by Airline

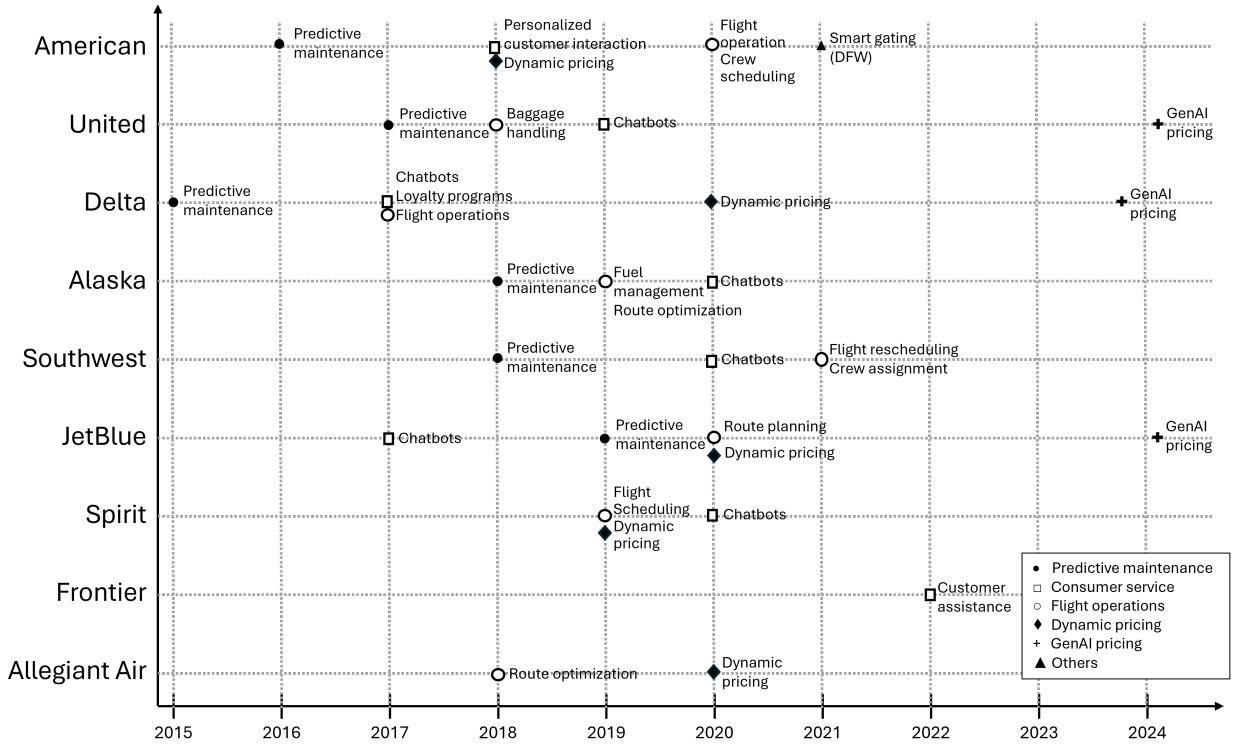
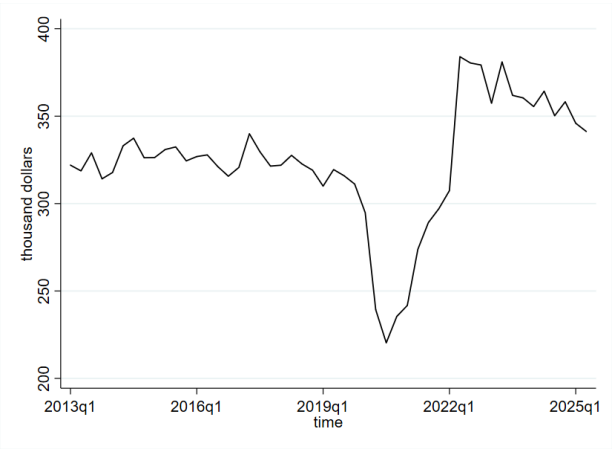
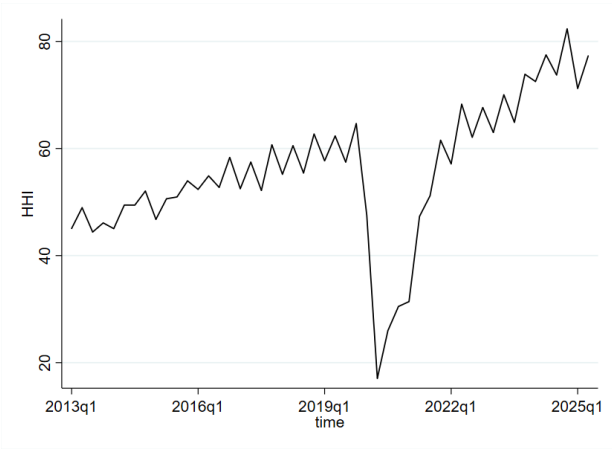


Figure A2: Evolution of Airfares, On-Time Performance, Market Concentration, and Air Quality

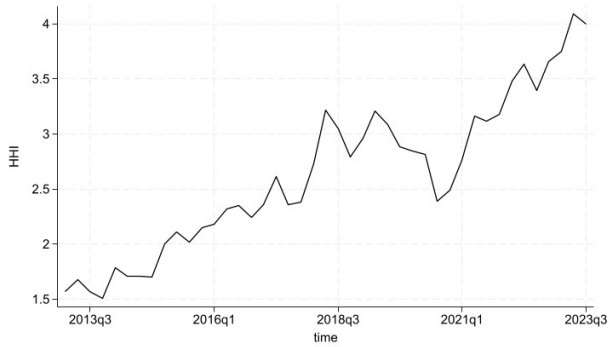
(a) Airfares



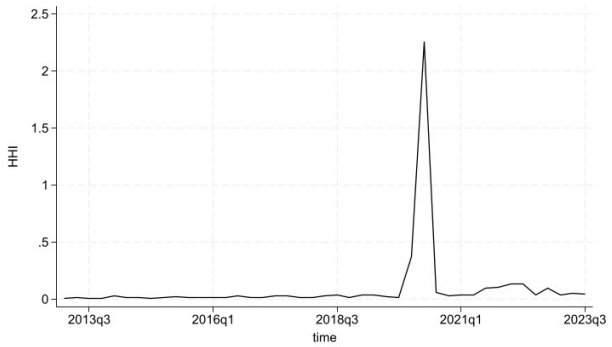
(b) HHI



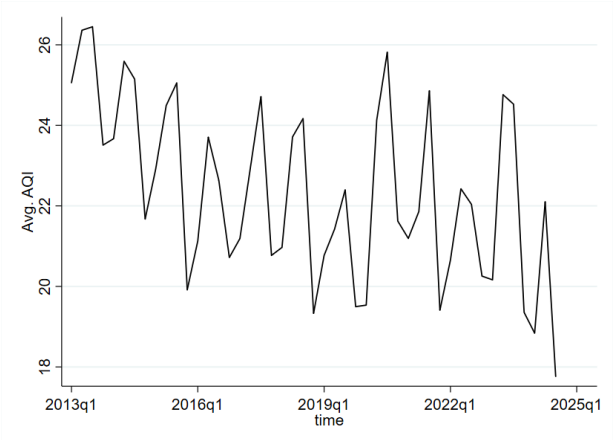
(c) Delay Time



(d) Cancellation (%)



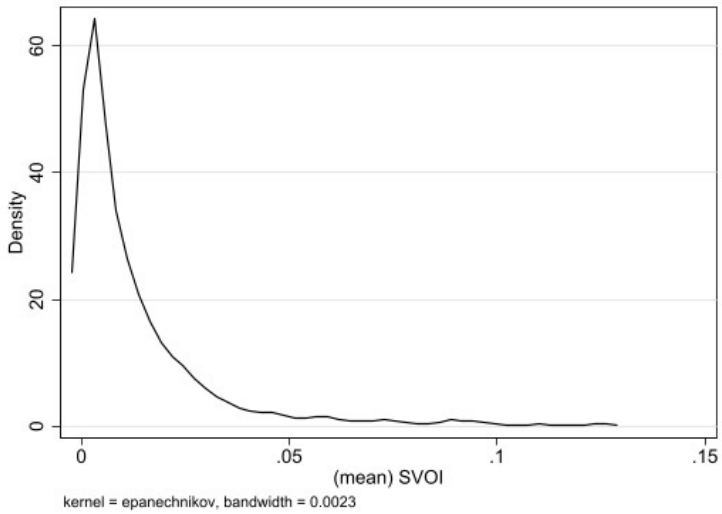
(e) AQI



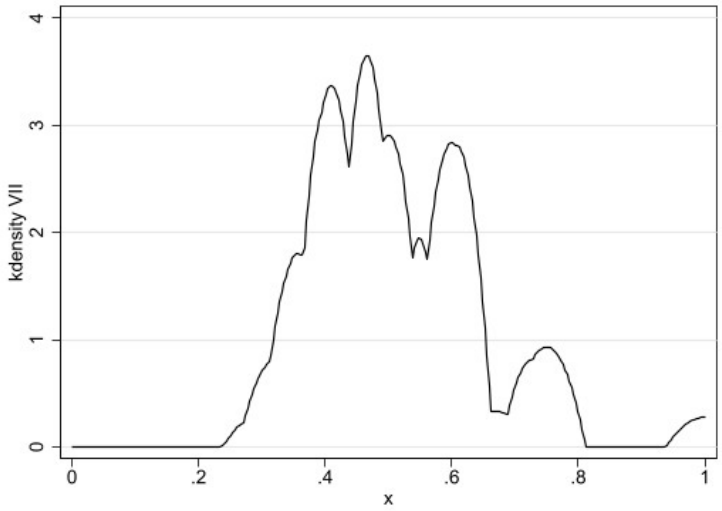
Note: Delay time and cancellation rate exclude delays due to weather-related reasons.

Figure A3: Distribution of GenAI Indices

(a) SVOI



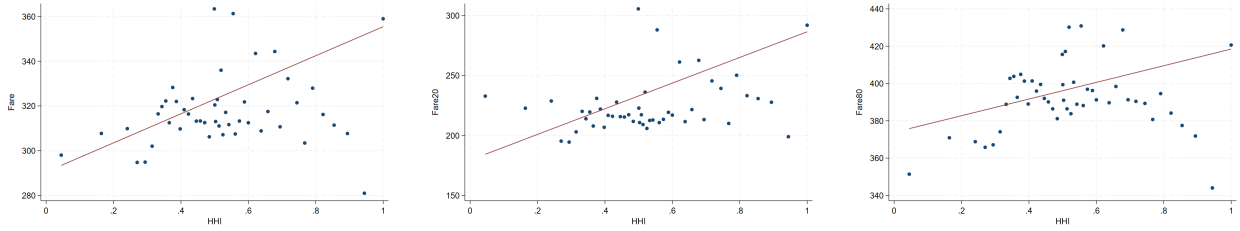
(b) VAII



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.

Figure A4: Market Concentration and Fares

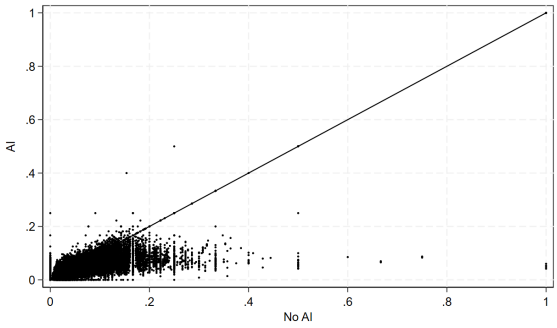
(a) lnFare (b) lnFare20 (c) lnFare80



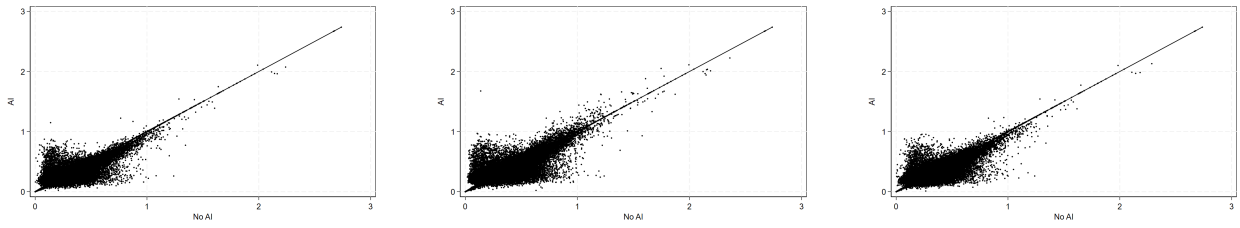
Note: Each panel shows a binned scatter plot of a fare measure against HHI. The HHI distribution is divided into equal-sized bins, and the mean fare within each bin is plotted. A fitted linear regression line is overlaid to illustrate the overall relationship between HHI and each fare measure.

Figure A5: Enterprise AI Adoption and Fare Distribution

(a) Share of low fares

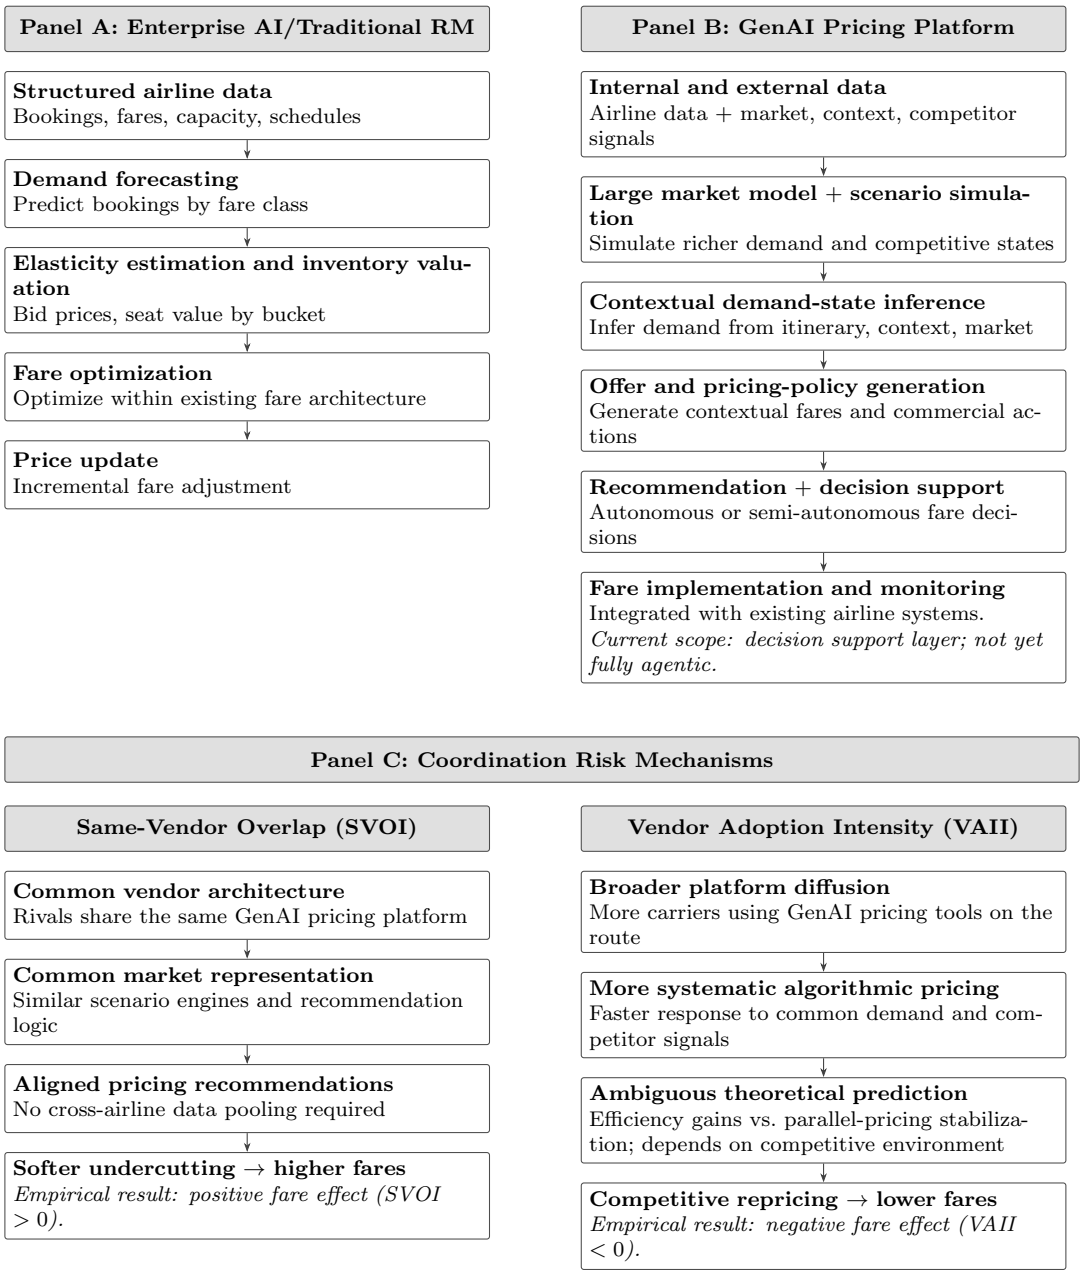


(b) Std. dev. of lnFare (c) Std. dev. of lnFare20 (d) Std. dev. of lnFare80



Note: Panel (a) compares the route-level share of low fares between periods when no airline had adopted enterprise AI and periods when all airlines on the route had adopted enterprise AI. Low fares are defined as fares below the route-specific 10th percentile computed over the full sample period. Panels (b)–(d) compare the standard deviation of route-level passenger-weighted fare measures between no-AI and all-AI periods. The solid line is the 45-degree line.

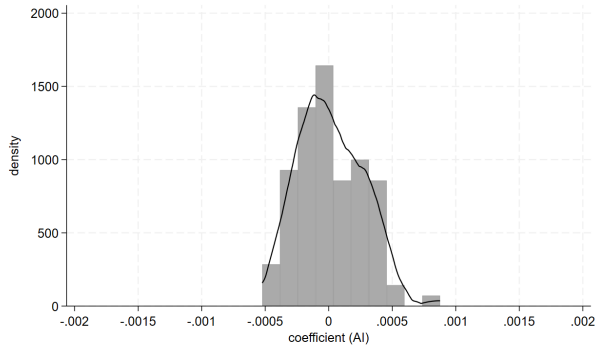
Figure A6: Enterprise AI, GenAI Pricing Platforms, and Coordination Channels in Airline Pricing



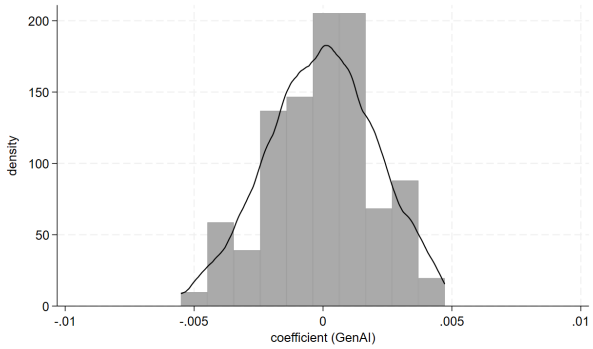
Note: Panel A depicts the enterprise AI forecast-and-optimize pipeline, in which structured airline data feed into demand forecasting, inventory valuation, fare optimization, and price updates. Panel B depicts the GenAI-enhanced simulate-and-recommend pipeline, in which internal airline data and external market signals are used to simulate market states and generate pricing recommendations within existing revenue management systems. Panel C summarizes the two coordination-risk channels examined in the empirical analysis. SVOI is the same-vendor overlap index, defined as the fraction of weighted airline pairs sharing the same GenAI pricing vendor on a route. VAII is the vendor adoption intensity index, defined as the share of airlines on a route using any third-party GenAI pricing vendor. The results from Table 3 indicate that SVOI is associated with higher fares, consistent with a common-vendor alignment channel, while VAII is associated with lower fares, consistent with a competitive-discipline channel from broader platform diffusion. See Sections 2.3, 3.2, 5.4, and 7.2.

Figure A7: Robustness Check: Placebo Test

(a) Placebo Test for Enterprise AI

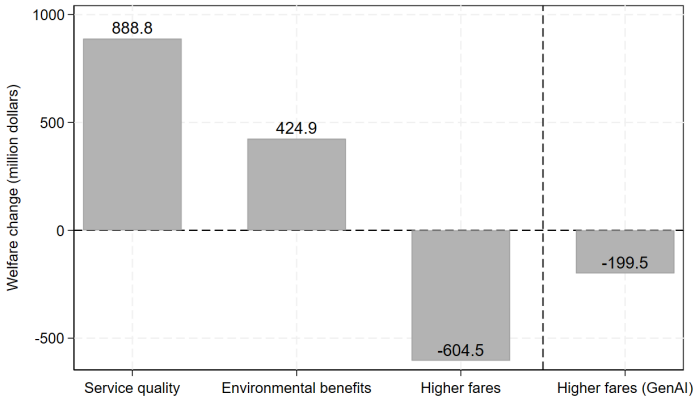


(b) Placebo Test for GenAI



Note: These figures report the distributions of estimated coefficients from separate placebo tests based on randomly reassigned enterprise AI adoption and GenAI platform use, respectively.

Figure A8: Welfare Impacts of AI Adoption



Note: The figure decomposes the consumer welfare impact of enterprise AI adoption into service-quality gains, environmental benefits, and fare-induced losses, expressed in millions of dollars. Welfare changes are computed using a nested logit demand model following Berry (1994), with utility depending on airfares, delay time, and AQI, and with nests defined by major airlines and LCCs. The enterprise-AI counterfactual removes AI adoption by reducing average fares by 1.01%, increasing delay time by 0.39 minutes, and increasing AQI by 2.16 units, based on the DID estimates in Table 1. Each component is expressed as compensating variation by dividing the utility difference by the estimated marginal utility of income. The GenAI counterfactual removes the estimated GenAI fare effect by reducing observed post-adoption fares by 5.13%. The corresponding compensating variation is calculated using data from 2013Q1–2024Q4, the period for which county-level AQI data are available.

Table A1: Enterprise AI vs. GenAI Pricing: Conceptual Distinctions and Empirical Measures

Dimension	**Enterprise AI**	**GenAI**
Core function	*Forecast-and-optimize.* Predict demand, value inventory, and optimize fares within the existing RM architecture.	*Simulate-and-recommend.* Simulate market states and generate contextual offers and pricing recommendations.
Role in pricing	Demand prediction, elasticity estimation, inventory valuation, fare optimization.	Market simulation, contextual offer generation, pricing-policy recommendation, decision support.
Data inputs	Booking histories, demand forecasts, competitor fares, operational records.	Internal airline data plus real-time demand, competitor prices, macroeconomic indicators, operational capacity, and market context.
Pricing mechanism	Optimizes fares within the existing RM architecture; improves forecast-based price and inventory decisions.	Generates pricing recommendations using market and contextual models; supports more automated decision support.
Implementation path	Hybrid use of internal systems, legacy RM workflows, vendor software, cloud infrastructure, and specialized AI modules.	Incumbent vendors such as PROS add contextual pricing, AI agents, and workflow automation; GenAI-focused vendors such as Fetcherr use market-simulation models to generate pricing recommendations.
Coordination channel	Faster detection of rivals' fare changes and faster response; reduces the short-run gain from undercutting in repeated pricing.	*Common-vendor alignment* (SVOI): shared architecture, market representation, and recommendation framework can soften undercutting. *Platform diffusion* (VAII): wider use of GenAI-enhanced pricing tools across carriers can increase responsiveness; fare effect is theoretically ambiguous.
Empirical measure	Enterprise AI adoption.	GenAI rollout; SVOI; VAII.

Note: SVOI = same-vendor overlap index; VAII = vendor adoption intensity index; RM = revenue management. The table distinguishes enterprise AI from GenAI-enhanced pricing platforms based on their primary role in the pricing process, not on a strict ownership or vendor classification. PROS is classified as GenAI in our empirical analysis when its pricing tools incorporate contextual pricing, AI agents, and workflow automation, while Fetcherr represents a GenAI-focused pricing vendor built around market-simulation models.

Table A2: Summary Statistics

	Enterprise AI Adoption			GenAI Platform Use		
	Full sample (1)	Not Adopted (2)	Adopted (3)	Full sample (4)	Not Adopted (5)	Adopted (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)			
Observations	3,697,957	1,055,346	2,642,611	5,488,575	5,079,504	409,071

Note: Means are reported at the airline–route–quarter level, with standard deviations in parentheses. Columns 1–3 use the enterprise AI sample from 2013Q1–2023Q3, and columns 4–6 use the GenAI sample from 2013Q1–2025Q2. Columns labeled *Not Adopted* and *Adopted* split observations according to whether the carrier had adopted the relevant technology by period t. SVOI and VAII are route–time measures of vendor commonality and platform diffusion, as defined in Section 3.2. Delay time and cancellation rates exclude weather-related disruptions.

Table A3: Impact of Enterprise AI Adoption (First Stage)

	(1)
Lag lnFare	0.934
	(1.476)
Lag Admin. Expense	0.019**
	(0.008)
Lag Res. Expense	-0.059**
	(0.028)
Lag Line Svc. Expense	0.005
	(0.033)
Lag #Routes	385.390*
	(232.773)
Lag #Tickets Sold	-0.000
	(0.000)
Lag BLP AI	8.344***
	(1.686)
LR χ^2	313.92
Observations	366

Note: This table reports the first-stage estimates for Table 1. 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. Each expense is calculated by multiplying annual expense and ticket shares. *Tickets Sold* denotes the total number of tickets sold by an airline in each year-quarter, and *BLP AI* is the average number of competitors who adopted enterprise AI. *** 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 A9. *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 ticket shares. *Origin SVOI* and *Destination SVOI* are average SVOI at the origin and destination airports, respectively. The sample includes Delta, United, American, Frontier, Southwest, and JetBlue. See Table 1 for other notes. *** p<0.01, ** p<0.05, * p<0.1.

Table A5: Robustness Check: Impact of Enterprise AI Adoption and GenAI Platform Use on Airfare Distribution (Interactive Factor Model)

	Enterprise AI			GenAI Platform Use		
	(1) lnFare	(2) lnFare20	(3) lnFare80	(4) lnFare	(5) lnFare20	(6) lnFare80
AI Adoption	0.018***	0.025***	0.016***			
	(0.001)	(0.001)	(0.001)			
GenAI				0.021***	0.007***	0.001
				(0.001)	(0.002)	(0.001)
Observations	623,448	623,448	623,448	671,111	671,111	671,111

Note: Sample periods are 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 10,000 passengers for columns 1–3 and more than 5,000 passengers for columns 4–6. All specifications are estimated using the interactive fixed effects model of Bai (2009), which allows latent common factors to have heterogeneous effects across units over time. All specifications include the same set of controls, airline fixed effects, route fixed effects, and time fixed effects as in Table 1. Standard errors in parentheses are clustered by route. *** p<0.01, ** p<0.05, * p<0.1.

Table A6: 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 3. *** p<0.01, ** p<0.05, * p<0.1.

Table A7: 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 periods are 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 1 for other notes. *** p<0.01, ** p<0.05, * p<0.1.

Table A8: 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. See Table 1 for other notes. *** p<0.01, ** p<0.05, * p<0.1.

Table A9: Robustness Checks: Impact of GenAI Platform Use and Vendor Sharing on Airfare Distribution

	(1) lnFare	(2) lnFare20	(3) lnFare80
Panel A. Alternative Fixed Effects			
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
Time FE	Y	Y	Y
Panel B. Control Function			
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

Note: Sample period is 2013Q1–2025Q2. The sample includes Delta, United, American, Frontier, Southwest, and JetBlue. All specifications include the same set of controls as in Table 1. Panel A replaces the separate airline and route fixed effects with airline-route fixed effects and includes time fixed effects. Panel B implements a control function approach, includes airline fixed effects, route fixed effects, and time fixed effects, and controls for residuals from the first-stage regressions for SVOI and VAII; the first-stage results are reported in Appendix Table A4. Standard errors in parentheses are clustered by route. *** p<0.01, ** p<0.05, * p<0.1.

Table A10: Robustness Check: Impact of Enterprise AI Adoption and GenAI Platform Use (Route-Level Estimates)

	Enterprise AI			GenAI Platform Use		
	(1) lnFare	(2) lnFare20	(3) lnFare80	(4) lnFare	(5) lnFare20	(6) 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)			
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	1,268,530	1,268,530	1,478,449	1,470,572	1,470,572	1,470,572

Note: Sample periods are 2013Q1–2023Q3 for columns 1–3 and 2013Q1–2025Q2 for columns 4–6. *AI Share* is the fraction of airlines that adopted enterprise AI in each route-time. *GenAI Share* is the fraction of airlines that adopted GenAI pricing in each route-time. All specifications include route-level controls, route fixed effects, and time fixed effects. Control variables are constructed as route-level averages of the corresponding variables. Instruments for lagged HHI, MMC, AI Share×HHI, and GenAI Share×HHI are constructed by interacting administrative expenses, reservation expenses, and line service expenses with the total number of tickets sold on each route. Standard errors in parentheses are clustered by airline. *** p<0.01, ** p<0.05, * p<0.1.

Table A11: Decomposing Airfare Effects: Quality and Capacity Channels

	(1) Markup	(2) lnFare	(3) Fare Effect	(4) lnDepartures	(5) lnSeats
AI Adoption	0.601***	0.010***		0.298***	0.294***
	(0.052)	(0.001)		(0.038)	(0.040)
Lag $\widehat{\text{HHI}}$		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	2,682,102	3,304,316	126,946	233,958	233,947

Note: *Markup* is calculated using the nested logit model based on Berry (1994). Details on the model are provided in Section 8.1. *lnDepartures* is the log of total departures performed, and *lnSeats* is the log of total seats derived from T-100 Segment data. All specifications include the same set of controls, airline fixed effects, and route fixed effects as in Table 1. Columns 1, 2, 4, and 5 also include time fixed effects; column 3 does not. 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 A14: Conduct Parameter Estimate

	(1) Markup
AI $\times$ μ^{pre}	0.121***
	(0.006)
μ^{pre}	0.176***
	(0.006)
Observations	2,492,473
Controls	Y
Airline $\times$ Time FE	Y
Route FE	Y
Time FE	Y

Note: Sample period is 2013Q1–2023Q3. The dependent variable is the cost-based actual markup, m_{ijt}. AI denotes the enterprise AI adoption indicator, and μ^{pre} denotes the pre-adoption average Bertrand markup for each airline-route pair. The coefficient on $AI \times \mu^{\text{pre}}$ is the conduct index θ, which measures whether enterprise AI adoption is associated with larger cost-based markup increases on airline-routes with greater pre-adoption Bertrand-implied markups. Control variables include average itinerary distance (measured by DB1B Market Distance) and its quadratic term, lagged origin and destination network size, the lagged number of tickets sold. Standard errors in parentheses are clustered by route. *** p<0.01, ** p<0.05, * p<0.1.

Table A12: 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
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

Note: Sample periods are 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. All specifications include the same set of controls, airline fixed effects, route fixed effects, and time fixed effects as in Table 1. Standard errors in parentheses are clustered by route. *** p<0.01, ** p<0.05, * p<0.1.

Table A13: 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

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, at which the implied net effect of GenAI on log average fare is positive ($-0.017 + 0.015 \times 3.76 = 0.039$). Thus, the overall effect of GenAI is positive at the sample mean of baseline scale. All specifications include the same set of controls, airline fixed effects, route fixed effects, and time fixed effects as in Table 3. Standard errors in parentheses are clustered by route. *** p<0.01, ** p<0.05, * p<0.1.