Are Auctions an Upgrade to Revenue Management? Evidence from the Airline Industry*

Alex Marsh† Garrett Scott‡ Jonathan W. Williams§

Last Updated: December 17, 2025

Abstract

We develop a model to study the profitability of complementing traditional revenue-management practices for selling vertically-differentiated goods with auctions and fixed-price sales to allocate upgrades. The effect on profits is ambiguous because strategic consumers can exploit the firm's lack of commitment by opting for the non-premium option and increasing the probability that upgrades are awarded in equilibrium. The magnitude depends on properties of the demand process and the level of awareness among consumers regarding the upgrade opportunities. To measure this trade-off in practice, we use data from a major airline that includes information on ticket sales, aircraft inventory, and awarded upgrades. For identification purposes, we exploit an experiment that introduced a pre-travel email to increase consumer awareness of the upgrade opportunities. We find that the notification effort doubled auction participation and bid revenue, but decreased overall revenue. However, consistent with the model's predictions regarding the heterogeneity of the effect on total revenue, we find a negative effect in business markets and positive effect in leisure markets. Taken together, the model and empirical findings highlight the challenges with integrating siloed algorithmic decision-making processes in settings with equilibrium effects.

Keywords: Auctions, Upgrades, Vertical Differentiation, Airlines, Price Discrimination

JEL Codes: L11, L21, L93, D44

*We thank the North American airline that provided the data used in this paper. We are particularly grateful to seminar and conference participants for useful comments, and Felipe Barbieri, Ben Rosa, Timothy Hubbard, and Geumbi Park for detailed discussions. NSF Grants SES 2417692 and 2417693 generously supported this research.

†Bates White Economic Consulting, **alex.marsh@bateswhite.com**

‡Department of Economics, University of Mississippi, **gescott1@olemiss.edu**

§Department of Economics, University of North Carolina at Chapel Hill, **jonwms@unc.edu**

1 Introduction

To further their profit-maximization objective, firms may combine sophisticated strategies to allocate differentiated goods to consumers. These hybrid approaches are prevalent in markets where fixed inventories of vertically-differentiated goods or services are sold to sequentially arriving customers.[1] In the airline industry, a common strategy is to complement revenue-management systems, which dynamically adjust posted prices to screen passengers between aircraft cabins, with auctions and fixed-price sales to allocate upgrades. In this paper, we develop a model to highlight the trade-offs associated with the introduction of these upgrade processes, and we measure the effect on consumer behavior and profits using data from an experiment by a major airline.

Our model considers a monopolistic airline that uses dynamic pricing to sell seats in two vertically-differentiated cabins before a fixed departure date, while offering opportunities for upgrades through an auction and fixed-price sales. Demand is characterized by a stochastic sequence of short-lived consumers with heterogeneous preferences who arrive to the market and decide whether to purchase a premium or economy seat, or nothing. A subset of the arrivals are informed about the upgrade processes and consider the value of bidding in the auction and the possibility of purchasing an upgrade at check-in when making their decision. Consistent with the practices of most airlines, see Hortaçsu et al. [2023] for a discussion of organizational siloing, the airline's pricing and upgrade decisions are made by separate teams guided by different algorithmic policies.[2]

We show that the pricing and upgrade teams' competing interests in the presence of strategic consumers can create circumstances that negatively impact profits. Specifically, the auction seeks to maximize revenue conditional on the number of seats remaining at the time of the auction, while fixed-price upgrades at check-in are offered contingent on availability. By opting not to buy premium seats outright, consumers increase the probability of winning the auction and the odds of seats being available at check-in for purchase because the airline cannot commit to withholding upgrades. This can reduce revenue directly by decreasing premium seat sales and indirectly via dynamic adjustments to fares that do not anticipate upgrades. Taken together, if upgrade revenue does not offset losses from outright sales, the airline's revenue can decrease. The sign and magnitude of the effect depend on the demand process. In particular, markets with more heterogeneity in valuations and greater changes in the composition of demand as departure approaches are more negatively impacted by increased awareness and participation in the upgrade processes.

To measure the impact of upgrade processes implemented in a manner consistent with the model, we use a recent and novel 9-month panel of data from a large North American airline.[3] The sample is drawn from 130 origin and destination (O&D) pairs (i.e., markets) and includes information on every ticket purchase, daily aircraft inventory records from each flight leading up

[1] Airline example: www.qantas.com/example. Hotel example: www.thepointsguy.com/hyatt.

[2] Heuristic approximations, or breaking apart larger optimization problems, is common for firms solving complex decisions with limited resources (e.g., Radner [1993]), including intractable network routing and pricing decisions in the airline industry (e.g., Barnhart and Sheffi [1993] and Barnhart et al. [2003]).

[3] Our data-use agreement places limits on the disclosure of some information (e.g., timing of experiment) for anonymity purposes.

to departure, and information on upgrades awarded through the auction and check-in sales.

The basic descriptive statistics of our data are representative of the broader airline industry during our sample period. Across all operated flights, the ratio of economy-cabin seats to premium-cabin seats is approximately eleven to one. Both cabins average a load factor (i.e., fraction of seats filled at departure) of just above 80%, with the premium cabin being more variable. Premium fares average approximately \$379, while economy fares average \$184, demonstrating a strong preference for quality among some consumers.[4]

At the beginning of our panel, the airline introduced an upgrade auction to complement the existing revenue-management system and check-in upgrade process. Bids were restricted to be among a set of discrete values on a slider between a minimum and maximum, and the customer could choose their bid by adjusting the slider from an initial position. The upgrade decisions, i.e., auction and fixed-price sales, were made independently from the revenue-management pricing decisions with algorithmic guidance from software provided by a third-party vendor. These upgrade processes were an important part of the airline's seat allocation strategy. Approximately 27% of passengers in the premium cabin at departure were awarded an upgrade through the auction or at check-in. Accepted bids and check-in fees average \$97 and \$104, respectively.

During the panel, we also observe the airline's experiment with an email notification system intended to raise awareness about the opportunity to bid in the upgrade auction. In the context of the model, this corresponds to an increase in the fraction of arrivals in the demand process that are informed about the auction. The email notifications approximately double the number of placed and accepted bids, with little change to the average bid.[5]

We exploit the notification experiment with an event-study design to identify a causal effect of the increase in awareness of and participation in the upgrade processes on the airline's profits. Our regression analysis focuses on two common measures of profitability: revenue per available seat (RAS) and revenue per available seat mile (RASM). Regardless of the size of the event window (e.g., 30-, 60-, or 90-days) or revenue measure, we find that greater participation increases premium revenues while decreasing economy and overall revenues. We estimate the decrease in total RAS to be between \$1.95 and \$2.77, depending on the event window. To elucidate the magnitude of the effects, we use an instrumental-variable (IV) regression framework to measure the impact of each additional bid due to the email notifications. We find that an additional bid decreases total RAS by between \$2.12 and \$3.23. Our estimates of the impact on RASM have a similar consequence for profitability after adjusting for distance flown.

To test the predictions of the model regarding heterogeneity in the effect of increased participation in the upgrade processes on profits, we use a taxonomy of markets created by the airline. Specifically, the airline classifies *business* markets as exhibiting greater heterogeneity in valuations and a more dynamic composition of demand as departure approaches, compared to *leisure* mar-

[4]The airline decomposes connecting itinerary fares into flight-segment components, so these fares correspond to the price paid on a nonstop flight segment.

[5]This result is consistent with the model's assumption that the initial selection of bidders was driven largely by awareness.

kets. Repeating the event-study regression analysis separately on each group of markets, we find that both RAS and RASM decrease in business markets, but increase in leisure markets. This is consistent with the model's prediction of substantial heterogeneity in the impact on profits across markets due to differences in the primitives of the demand process.

Taken together, our model and empirical findings highlight challenges for practitioners when complementing revenue management with ex-post reallocation strategies like auctions. The primary challenge is that strategic consumers can exploit a siloed algorithmic decision-maker's lack of commitment to withhold upgrades. If the upgrade processes are not integrated and coordinated with pricing, consumers can adapt their behavior to create equilibrium conditions (e.g., unfavorable outright premium sales) that validate the belief that a less-expensive upgrade is likely to be available. This can alter the path of offered fares and create, rather than resolve, allocative inefficiencies (e.g., large number of upgraded consumers crowding out economy sales prior to departure). Further, any effort to measure the return from implementation should carefully account for the possibility of heterogeneity that may mask meaningful failures and successes.

Our paper contributes to the literature in operations and economics studying allocative mechanisms, particularly pricing and auctions. The literature on dynamic pricing has numerous notable theoretical (e.g., Stokey [1979], Gale and Holmes [1993], Dana [1999], Courty and Li [2000], Armstrong [2006]) and empirical (e.g., Graddy and Hall [2011], Sweeting [2010], Cho et al. [2018], Waisman [2021], Williams [2022], Escobari [2012]) contributions. The workhouse models in this literature have been enriched to include competition (e.g., Gallego and Hu [2014] and Betancourt et al. [2022]), discounting (e.g., Dilmé and Li [2018] and Dilmé and Garrett [2022]), and multiproduct firms (e.g., Aryal et al. [2023], Maglaras and Meissner [2006], Dong et al. [2009], Talluri and van Ryzin [2004]).[6] Notable contributions to the literature on the use of auctions to allocate multiple units of homogeneous and differentiated goods include Vulcano et al. [2002] and Gershkov and Moldovanu [2009], respectively. Regarding the airline industry, auctions have been explored as a way to allocate seats (e.g., Talluri and van Ryzin [1998]) and resolve over-bookings (e.g., Ely et al. [2017]).

Our work demonstrates the challenges of combining strategies to allocate goods, e.g., dynamic pricing and auctions. In our setting, the potential for negative profit outcomes arises because of endogenous entry into the auction, (e.g., Samuelson [1985], Levin and Smith [1994], Marmer et al. [2013], Roberts and Sweeting [2013], Gentry and Li [2014], Gentry et al. [2017]). Strategic consumers can exploit the state-contingent profit-maximization objective of the auction by selecting into the upgrade processes, thus undermining the screening intent of dynamically-adjusted prices. Similarly, Jerath et al. [2010] highlights the need for firm's strategies to account for the harm to profits due to anticipation of advantageous opportunities by consumers (e.g., last-minute discounts on fares to sell unused capacity). Shen and Su [2007] provide an excellent overview of issues related to consumer behavior in revenue-management models. Recent empirical research on the airline industry demonstrates that consumers are also strategic with respect to the timing of purchases

[6] Airlines also commonly use bundling to sell ancillary services like checked bags (e.g., Cui et al. [2018b]).

(e.g., Li et al. [2014] and Li et al. [2019]) and refundability (e.g., Lazarev [2013] and Scott [2024]).

Most closely related to our work is research on allocating upgrades. Gallego and Phillips [2004] study a model of "flexible products" used by firms to discriminate among customers based on their willingness to adapt or compromise (e.g., timing of flights), which can have benefits for capacity management. In a digital environment, it is convenient and straightforward to provide customers with opportunities to enrich their experience in real-time. Sheng et al. [2025] study this problem in the gaming industry and develop optimal strategies for selling "bonus actions" that provide a boost to help a player achieve an objective. Cui et al. [2018b] provide a rich and flexible modeling framework and characterize both equilibrium behavior of randomly-arriving consumers that anticipate the possibility of upgrades and the optimal pricing of the upgrade. Specific to airlines, Cui et al. [2018a] find that fare dispersion and revenues from the economy cabin increase after the introduction of an option to upgrade to premium seating. Favrizi [2024] develops a model to quantify the welfare consequences of airlines that set prices dynamically and offer fixed-price upgrades to myopic consumers who do not anticipate the possibility of an upgrade when making initial purchase decisions. Our model combines aspects from each these studies to include dynamic pricing, upgrade processes, and strategic consumers, while our empirical analysis offers insight into the ambiguity regarding the returns from a common form of implementation in the airline industry.

The remainder of the paper is as follows. Section 2 presents the theoretical model and numerical simulations to demonstrate the effect of increased awareness on revenues. Section 3 describes our data and provides descriptive statistics. Section 4 presents the event-study and IV regression frameworks and results, while Section 5 concludes.

2 Model

In this section, we develop a model to highlight the trade-offs associated with integrating upgrade processes into the revenue-management systems of firms selling vertically-differentiated goods to strategic consumers. We then conduct numerical simulations to assess how increased awareness of and participation in these upgrade processes influence consumer behavior and profitability, and to identify the potential sources of heterogeneity in these effects. The insights from the model and simulations guide our approach to the empirical analysis in Section 4.

2.1 Overview

Timing We consider a monopolistic airline that maximizes profit from a given flight by allocating premium (f) and economy (e) cabin seats to a stochastically arriving sequence of consumers with heterogeneous preferences. Time is discrete and indexed by $t \in \{1, ..., T, T+1\}$. A flight has initial capacity $\boldsymbol{k}_1 = (k_1^f, k_1^e)$, representing the total number of seats in the premium and economy cabins, respectively. Let $\boldsymbol{k}_t = (k_t^f, k_t^e)$ denote the remaining capacity in period t.

In each period $t \leq T$, the airline posts cabin-specific prices $\boldsymbol{p}_t = (p_t^f, p_t^e)$ and sets seat-release policies $\overline{\boldsymbol{q}}_t = (\overline{q}_t^f, \overline{q}_t^e)$, where $\overline{q}_t^f \leq k_t^f$ and $\overline{q}_t^e \leq k_t^e$, before the arrival of consumers. These release policies correspond to the maximum number of seats that can be sold in period t at prices $\boldsymbol{p}_t$.

Consumers who purchase a seat in the economy cabin in any period $t < \tilde{t}$ may submit a bid for an upgrade to the premium cabin, where $\tilde{t} < T + 1$ is the period in which bids are accepted or rejected. In $T + 1$, consumers check in for the flight, and the airline offers upgrades to consumers at a fixed fee if premium seats are available. The flight departs at the end of $T + 1$.

Consumer Arrivals and Preferences We assume that N_t consumers arrive in each period $t \leq T$, where N_t is Poisson distributed with mean λ_t. Consumers choose from the set $\{f, e, o\}$, where o represents the outside option, and cannot delay their decision. A fraction $\alpha \in [0, 1]$ of the consumers are aware of the upgrade processes.

The utility of a consumer i that arrives in period t and is unaware of the upgrade processes is

$$u_{it}^m \equiv \nu_i \xi_i^m - p_t^m,$$

for each of the choices $m \in \{f, e, o\}$. ν_i represents consumer i's willingness to pay for the flight, ξ_i^m captures consumer i's perceived quality of option m, and p_t^m equals the price of option m. The quality of the outside option, ξ_i^o, and its price, p^o, are normalized to 0 for all i. Similarly, we normalize the quality of the economy cabin to $\xi_i^e = 1$ for all i, which permits $\xi_i^f = \xi_i$ to be interpreted as the premium cabin's quality relative to economy.

If a consumer is aware of the upgrade processes, the utilities of choices f and o remain unchanged, but the economy ticket gains the option value associated with the probability of being upgraded and traveling in the premium cabin. In this case, consumer i's expected utility from choice e is

$$\mathcal{U}_{it}^e = \nu_i - p_t^e + \varrho_{it}^* \big(\nu_i(\xi_i - 1) - b_{it}^*\big) + \varphi_t^* \big(1 - \varrho_{it}^*\big) \max\big\{0, \nu_i(\xi_i - 1) - r\big\}.$$

The utility of the economy cabin, $\nu_i - p_t^e$, is guaranteed to consumer i with the purchase of the economy fare. However, there are two channels through which the consumer can be upgraded.

Let ϱ_{it}^* be consumer i's beliefs about the probability of winning an upgrade through the auction given a bid of b_{it}^* from among a discrete set of permissible values, and φ_t^* be the belief about the availability of an upgrade during check-in at price r. Then, the option value created for the consumer by the upgrade processes equals $\varrho_{it}^* \big(\nu_i(\xi_i - 1) - b_{it}^*\big) + \varphi_t^* \big(1 - \varrho_{it}^*\big) \max\big\{0, \nu_i(\xi_i - 1) - r\big\}$. The first term captures the expected surplus from winning the auction, and the second reflects the expected surplus from a check-in upgrade conditional on not winning the auction. We discuss the calculation of equilibrium beliefs, ϱ_{it}^* and φ_t^*, below.

To capture heterogeneity in preferences among consumers, we assume that there are two types: $\mathcal{L}$ and $\mathcal{B}$. Upon arrival, type $\mathcal{L}$ ($\mathcal{B}$) consumers draw their realization of ν from a type-specific exponential distribution with mean $\lambda_\nu^{\mathcal{L}}$ ($\lambda_\nu^{\mathcal{B}}$), while ξ is equal to one plus a draw from a type-specific exponential distribution with mean $\lambda_\xi^{\mathcal{L}}$ ($\lambda_\xi^{\mathcal{B}}$). We assume that the probability a consumer is type $\mathcal{B}$ equals zero in period $t = 1$ and increases each period by $\Delta^{\mathcal{B}}$, such that the probability in period t equals $\min\left(\Delta^{\mathcal{B}}(t-1), 1\right)$.

Pricing and Upgrade Policies As is common in the airline industry, decisions regarding pricing and upgrades are made by two separate teams that are not integrated. In particular, pricing policies are not anticipatory of the option value created by the upgrade processes, and the decision to award upgrades does not internalize the equilibrium effect on consumers' decision-making that impacts the distribution of states at the time of the auction. The siloed nature of the decisions substantially simplifies solving for the optimal policies of the airline for a demand process given by $\left(\alpha, \lambda_{\nu}^{\mathcal{L}}, \lambda_{\nu}^{\mathcal{B}}, \lambda_{\xi}^{\mathcal{L}}, \lambda_{\xi}^{\mathcal{B}}, \Delta^{\mathcal{B}}\right)$.

First, consider the pricing decision. In each $t \leq T$, the airline sets prices $\boldsymbol{p}_t = (p_t^f, p_t^e)$ and releases a number of seats for sale $\overline{\boldsymbol{q}}_t = (\overline{q}_t^f, \overline{q}_t^e)$ prior to the arrival of consumers. Let $\mathbb{E}_t[Q^m(\boldsymbol{p}, \overline{\boldsymbol{q}})]$ be the expected quantity demanded in cabin m at time t given prices $\boldsymbol{p}$ and released seats $\overline{\boldsymbol{q}}$. Then, the airline's expected per-period revenues can be written $\mathbb{E}_t[R(\boldsymbol{p}, \overline{\boldsymbol{q}})] = p^f \mathbb{E}_t[Q^f(\boldsymbol{p}, \overline{\boldsymbol{q}})] + p^e \mathbb{E}_t[Q^e(\boldsymbol{p}, \overline{\boldsymbol{q}})]$.

The airline faces a constant marginal "peanut" cost from servicing a passenger in cabin m equal to c^m, with $c^e \leq c^f$ and $\boldsymbol{c} = (c^f, c^e)$. Total peanut costs equal $C(\boldsymbol{k}) = c^f(k_1^f - k^f) + c^e(k_1^e - k^e)$, where $(k_1^f - k^f)$ and $(k_1^e - k^e)$ are the number of occupied seats in the premium and economy cabins, respectively.

In period T, the pricing team's optimal policies solve

$$V_T(\boldsymbol{k}) = \max_{\substack{\overline{\boldsymbol{q}} \in \mathcal{K}_T \\ \boldsymbol{p} \in \mathbb{R}_+^2}} \mathbb{E}_T\left[R\left(\boldsymbol{p}, \overline{\boldsymbol{q}}\right)\right] - \int_{\boldsymbol{k}' \in \mathcal{K}} C(\boldsymbol{k}') dH_T(\boldsymbol{k}' | \boldsymbol{k}, \boldsymbol{p}), \tag{1}$$

given that revenue and costs from check-in upgrades are not considered (i.e., the pricing team perceives seats to have zero value after period T, such that $V_{T+1}(\boldsymbol{k}) = 0$). $\mathcal{K}_T = \{0, 1, ..., k_T^f\} \times \{0, 1, ..., k_T^e\}$ denotes the subset of the state space available for setting seat-release policies, and $H_T(\boldsymbol{k}' | \boldsymbol{k}, \boldsymbol{p})$ is the distribution of next period's state conditional upon the state and policies in period T.

In period $t < T$, the airline's optimal policies solve

$$V_t(\boldsymbol{k}) = \max_{\substack{\overline{\boldsymbol{q}} \in \mathcal{K}_t \\ \boldsymbol{p} \in \mathbb{R}_+^2}} \mathbb{E}_t\left[R\left(\boldsymbol{p}, \overline{\boldsymbol{q}}\right)\right] + \int_{\boldsymbol{k}' \in \mathcal{K}} V_{t+1}(\boldsymbol{k}') dH_t(\boldsymbol{k}' | \boldsymbol{k}, \boldsymbol{p}). \tag{2}$$

As with check-in upgrades, the pricing team's optimal policies for $t \neq T$ are without consideration of the auction in period $\tilde{t}$. This finite-horizon dynamic program can be solved recursively to yield the pricing team's policy functions $\left(\text{i.e., } \boldsymbol{p}_t(\boldsymbol{k}) \text{ and } \overline{\boldsymbol{q}}_t(\boldsymbol{k})\ \forall t\right)$ and value function $\left(\text{i.e., } V_t(\boldsymbol{k})\ \forall t\right)$.

Now, consider the upgrade team's decisions at check-in and the auction. Given that a seat has zero value at departure, the optimal decision at check-in is to accept all customers willing to pay the upgrade price of r up to remaining premium capacity after period T sales are complete. The profit from n check-in upgrades equals $\left(r - (c^f - c^e)\right) n$.

Let $U_{T+1}(\boldsymbol{k}) = \left(r - (c^f - c^e)\right) \mathbb{E}_{T+1}(n)$ denote the expected profit from check-in upgrades with capacity $\boldsymbol{k}$ remaining at time $T+1$. The expectation is taken over the valuations of consumers that purchased an economy seat and were not awarded an upgrade through the auction.

The decision to accept a bid must balance bid revenue net of additional peanut costs against the opportunity cost of reallocating passengers between cabins, which captures any impact on expected revenue from outright sales of seats and fixed-price upgrade sales at check-in. Denote the value function at period T, inclusive of outright sales and upgrade opportunities, as

$$U_T(\boldsymbol{k}) = \max_{\substack{\overline{\boldsymbol{q}} \in \mathcal{K}_T \\ \boldsymbol{p} \in \mathbb{R}^2_+}} \mathbb{E}_T\big[R\big(\boldsymbol{p}, \overline{\boldsymbol{q}}\big)\big] + \int_{\boldsymbol{k}' \in \mathcal{K}} \Big(U_{T+1}(\boldsymbol{k}') - C(\boldsymbol{k}')\Big) dH_T(\boldsymbol{k}'|\boldsymbol{k}, \boldsymbol{p}). \tag{3}$$

Continuing with this recursive representation, the value function at the time of the auction in period $\tilde{t}$ equals

$$U_{\tilde{t}}(\boldsymbol{k}) = \max_{\substack{\overline{\boldsymbol{q}} \in \mathcal{K}_T \\ \boldsymbol{p} \in \mathbb{R}^2_+}} \mathbb{E}_{\tilde{t}}\big[R\big(\boldsymbol{p}, \overline{\boldsymbol{q}}\big)\big] + \int_{\boldsymbol{k}' \in \mathcal{K}} U_{\tilde{t}+1}(\boldsymbol{k}') dH_{\tilde{t}}(\boldsymbol{k}'|\boldsymbol{k}, \boldsymbol{p}). \tag{4}$$

Variation in $U_{\tilde{t}}(\boldsymbol{k})$ across states in period $\tilde{t}$ captures the value of unoccupied seats and opportunities for profitable reallocation from accepting bids.

The optimal bid acceptance policy is to accept bids until the cost associated with reallocation, i.e., the impact on expected future outright sales and check-in upgrades, is greater than the next highest bid. Let $\boldsymbol{i}^u = (-1, 1)$ be an operator that decreases premium capacity and increases economy capacity by one seat. Then, accepting n bids at time $\tilde{t}$ implies a change in the capacity state from $\boldsymbol{k}_{\tilde{t}}$ to $\boldsymbol{k}_{\tilde{t}} + n\boldsymbol{i}^{\boldsymbol{u}}$ at a cost equal to $U_{\tilde{t}}(\boldsymbol{k}_{\tilde{t}} + n\boldsymbol{i}^{\boldsymbol{u}}) - U_{\tilde{t}}(\boldsymbol{k}_{\tilde{t}})$. The marginal cost of the n^{th} upgrade, denoted $\Delta U_{\tilde{t}}(n, \boldsymbol{k})$, equals

$$\Delta U_{\tilde{t}}(n, \boldsymbol{k}) = \begin{cases} 0 & \text{if } n = 0 \\ U_{\tilde{t}}(\boldsymbol{k} + n\boldsymbol{i}^{\boldsymbol{u}}) - U_{\tilde{t}}(\boldsymbol{k} + (n-1)\boldsymbol{i}^{\boldsymbol{u}}) & \text{if } n \in \{1, 2, ..., k^f\} \\ \infty & \text{otherwise.} \end{cases}$$

Letting $b^{(n)}$ denote the n^{th} largest bid received, the optimal policy is to accept n bids if

$$\Delta U_{\tilde{t}}(n, \boldsymbol{k}) < b_{\tilde{t}}^{(n)} \quad \text{and} \quad \Delta U_{\tilde{t}}(n+1, \boldsymbol{k}) > b_{\tilde{t}}^{(n+1)}.$$

Together, the pricing, seat-release, and bid acceptance policies characterize the airline's optimal choices for a given demand process.

Equilibrium Given a demand process and cabin-specific costs, which are defined by the set of parameters $\Theta = \big[\big(\alpha, \lambda_\nu^{\mathcal{L}}, \lambda_\nu^{\mathcal{B}}, \lambda_\xi^{\mathcal{L}}, \lambda_\xi^{\mathcal{B}}, \Delta^{\mathcal{B}}\big), \big(c^f, c^e\big)\big]$, a solution to the model includes pricing, seat-release, and bid-acceptance policies for the airline and equilibrium beliefs for consumers. Solving the model proceeds in two steps. First, given the siloed nature of the decision making, the pricing and seat-release policies are solved ignoring the influence of the upgrade processes. As detailed in Aryal et al. [2023], this can be done using backward induction by solving a mixed-integer programming problem at each state. The optimal policies and associated value functions for the given demand and cost parameters (i.e., $\boldsymbol{p}_t(\boldsymbol{k};\Theta)$, $\overline{\boldsymbol{q}}_t(\boldsymbol{k};\Theta)$, and $V_t(\boldsymbol{k};\Theta)$ $\forall t$) can then be used to solve for equilibrium consumer beliefs and the value functions underlying bid-acceptance policies (i.e., $(\varrho^*_{it},\varphi^*_t)$ and $U_t(\boldsymbol{k};\Theta)$ $\forall t$, respectively) using the forward-simulation procedure of Marsh et al. [2024b]. The algorithm simulates consumers' choices given an initial set of beliefs along with an optimal response by the airline to those choices and calculates the resulting probabilities of being awarded an upgrade in the auction and at check-in. Those two steps are then repeated until beliefs converge to the win probabilities yielded by the airline's optimal policies.

2.2 Numerical Simulations

Figure 1: Entry and Bidding Behavior by Share of Aware Consumers

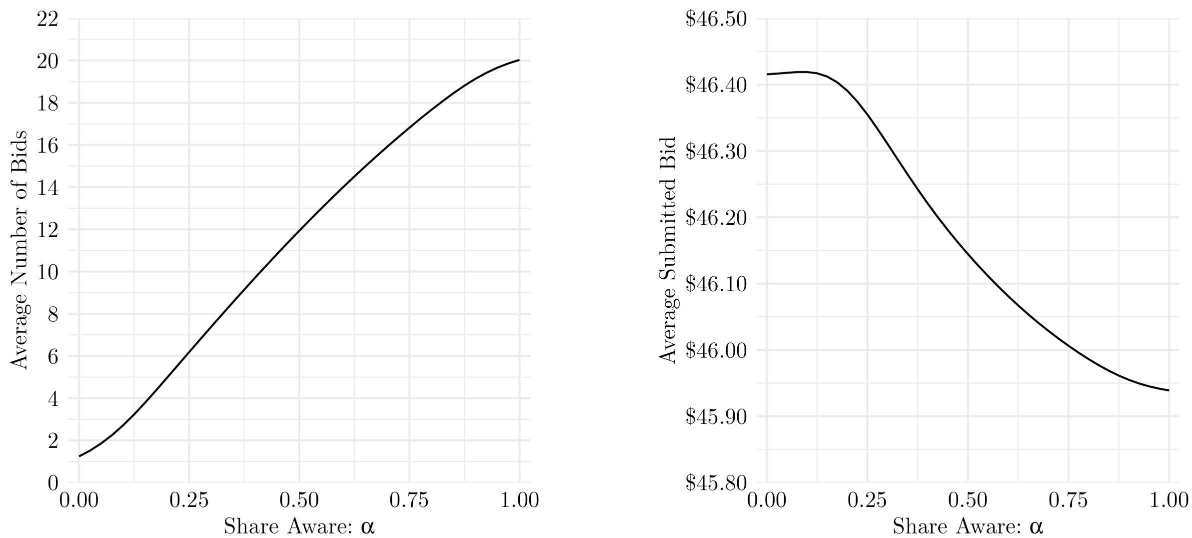

(a) Average Number of Bids

(b) Average Bid

Notes: Panel (a) presents the relationship between awareness and the average number of bids. Panel (b) presents the relationship between awareness and the average value of submitted bids. The calculations are performed using 75,000 simulations (i.e., flights) of the model's equilibrium and a Gaussian kernel is used to smooth simulation error.

To motivate and guide our empirical analysis in Section 4, we simulate the model to demonstrate how an increase in consumer awareness of upgrade processes, i.e., increase α, can affect airline profitability. Specifically, we seek to establish whether increased awareness can reduce profits, and whether this effect varies with the other parameters of the demand process that likely differ across markets.

In our model, the possibility of reduced profits arises because the airline's pricing and upgrade decisions are not fully integrated, which leads to a lack of commitment from the airline to withhold upgrade opportunities from consumers that would otherwise buy a premium seat outright. As the fraction of consumers that are aware of the upgrades increases, the screening intention of cabin-specific prices may be undermined, and profits can decrease if lost outright sales revenue is not offset by revenue from upgrades.

To demonstrate this possibility, we consider the following demand and cost parameters,

$$\Theta = \left[\left(\alpha \in [0,1], \lambda_\nu^{\mathcal{L}} = 11.74, \lambda_\nu^{\mathcal{B}} = 27.89, \lambda_\xi^{\mathcal{L}} = 1.15, \lambda_\xi^{\mathcal{B}} = 1.68, \Delta^{\mathcal{B}} = 0.067 \right), \left(c^f = 10, c^e = 5 \right) \right],$$

and solve the model for a grid of α values, $\alpha \in \{0, 0.025, ..., 0.975, 1\}$. For simplicity, we permit consumers to bid among ten equally spaced discrete values ranging from $b^1 = 44.25$ to $b^{10} = 88.50$, and set $r = 59$, $T = 15$, $\tilde{t} = 14$, and $(k_1^f, k_1^e) = (12, 162)$.[7]

For each value of Θ, we solve the model and simulate 75,000 flights using consumers' equilibrium beliefs and the airline's optimal pricing and check-in policies. This approach holds preferences fixed while measuring how efforts to increase consumer awareness affect profits.

Figure 2: Upgrade Revenues and Profits Per Flight by Share of Aware Consumers

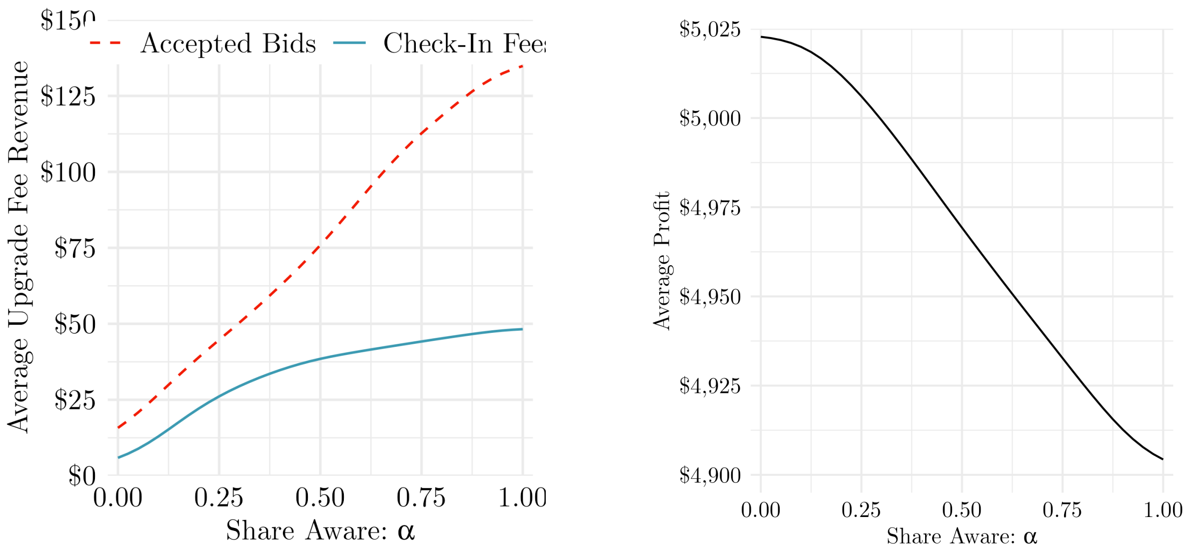

(a) Average Upgrade Revenue (b) Average Flight Profit

Notes: Panel (a) presents the relationship between awareness and average per-flight revenue from the auction and check-in upgrades. Panel (b) presents the relationship between awareness and the per-flight total profits, which includes fares and upgrade revenue. The calculations are performed using 75,000 simulations (i.e., flights) of the model's equilibrium and a Gaussian kernel is used to smooth simulation error.

Figure 1 shows how bidding behavior in the auction changes with α. In panel (a), the average number of bids per flight increases monotonically in α, consistent with increased consumer awareness of the upgrade processes. In panel (b), the average submitted bid decreases slightly as α increases. This can occur, despite greater competition among bidders, because participation is endogenous and prices adjust to the altered evolution of states that results from consumers changing their decisions.[8]

Figure 2 illustrates how greater consumer awareness affects profits. Panel (a) shows that average per-flight upgrade revenues from accepted bids and check-in sales increase monotonically in α.

[7] As discussed in Section 3, this is consistent with the discrete options given to consumers by the airline.

[8] See Li and Zheng [2009] for settings in which equilibrium bidding behavior becomes less aggressive with a greater number of potential bidders.

Figure 3: Awareness and Profitability: Comparative Statics for $\Delta^{\mathcal{B}}$

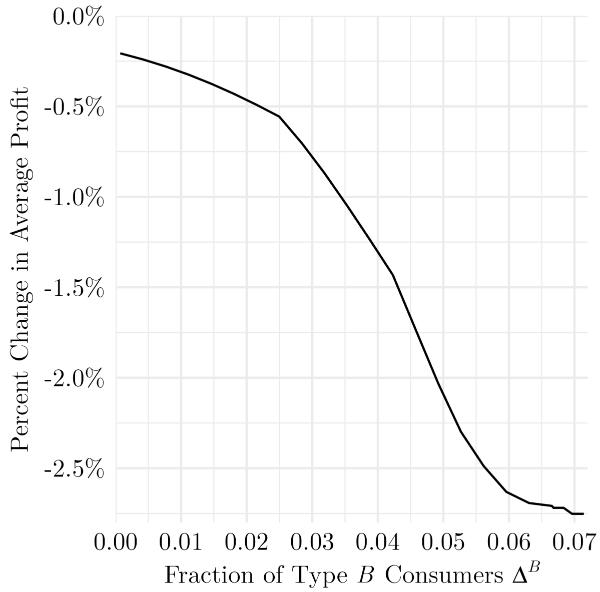

Notes: The figure plots the percentage change in average equilibrium profits when all consumers are aware of the upgrades for different values of the per-period increase in the probability of a business traveler changes, holding all other parameters constant. The calculations are performed using 75,000 simulations (i.e., flights) of the model's equilibrium and a Gaussian kernel is used to smooth simulation error.

Despite upgrade revenue increasing with α, panel (b) shows that average per-flight profit decreases. This is due to a larger share of premium seats being awarded through upgrades, rather than outright sales through the revenue-management system.

It is important to note that these findings come from a single specification of preferences, and that alternative specifications may yield different results. In the baseline specification, increasing α from zero (i.e., no customers aware) to one (i.e., full awareness) decreases profits by approximately 2.75%. To understand how changes in the demand process affect the results, we perturb each parameter individually and compare the impact on profits from greater awareness to this baseline loss of 2.75%.

For example, consider the rate of change in the proportion of business travelers, $\Delta^{\mathcal{B}}$. In the baseline specification, $\Delta^{\mathcal{B}} = \frac{1}{15}$, which is a rate that increases the fraction of business travelers from zero in period $t = 1$ to nearly one in period $T = 15$. Figure 3 shows the impact of increasing awareness from zero to one on profits for $\Delta^{\mathcal{B}} \in [0, \frac{1}{14}]$. We find that the percentage loss from increasing awareness is less in markets with a smaller rate of change of business consumers. This is because high-valuation business consumers are more likely to have paid an expensive premium fare but now choose an economy fare due to the option value associated with upgrades.

Figure 4 presents a similar comparative static exercise for the other parameters of the demand process, $\left(\lambda_{\nu}^{\mathcal{L}}, \lambda_{\nu}^{\mathcal{B}}, \lambda_{\xi}^{\mathcal{L}}, \lambda_{\xi}^{\mathcal{B}}\right)$, considering perturbations from the baseline specification. Panel (a) shows how perturbations of $\lambda_{\nu}^{\mathcal{L}}$ and $\lambda_{\nu}^{\mathcal{B}}$ affect profits relative to the baseline. The graph is centered at the baseline value for each parameter, i.e., a loss of approximately 2.75% from increasing awareness, and we consider deviations ranging from -20% to 20% for each parameter. We find greater values

Figure 4: Awareness and Profitability: Comparative Statistics for $\left(\lambda_\nu^{\mathcal{L}}, \lambda_\nu^{\mathcal{B}}, \lambda_\xi^{\mathcal{L}}, \lambda_\xi^{\mathcal{B}}\right)$

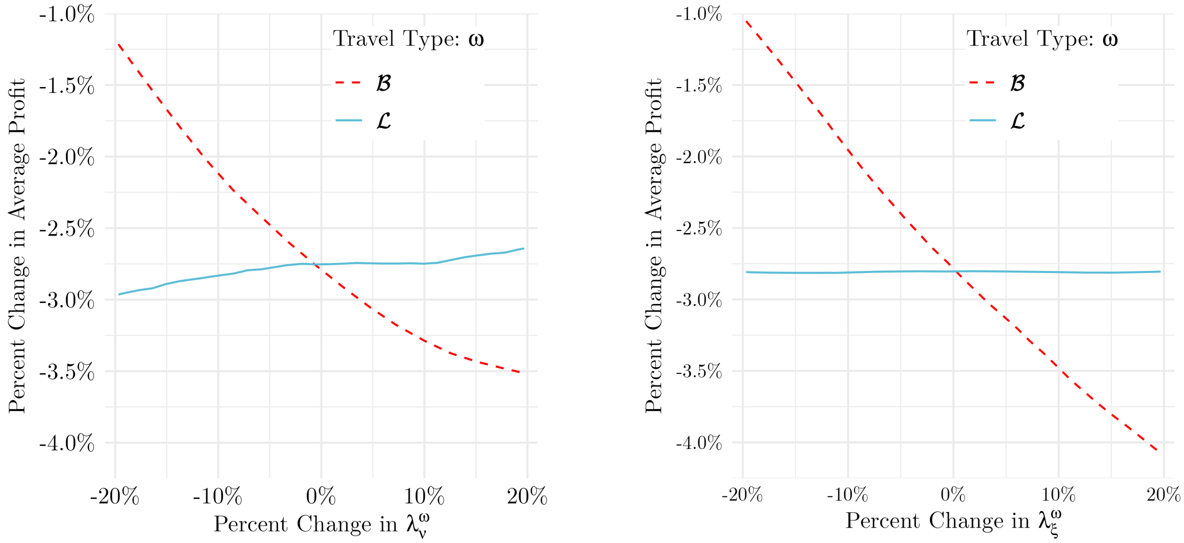

(a) WTP for Travel

(b) Preference for Quality

Notes: Panel (a) plots the percentage change in average equilibrium profits when all consumers become aware of the upgrades for different values of the willingness to pay of each type of customer is altered, i.e., $\lambda_\nu^{\mathcal{L}}$ and $\lambda_\nu^{\mathcal{B}}$, holding all other parameters constant. Panel (b) plots the same calculation for $\lambda_\xi^{\mathcal{L}}$ and $\lambda_\xi^{\mathcal{B}}$, holding all other parameters constant. The calculations are performed using 75,000 simulations (i.e., flights) of the model's equilibrium and a Gaussian kernel is used to smooth simulation error.

of $\lambda_\nu^{\mathcal{L}}$ correspond to a slightly less negative impact from increasing awareness, while greater values of $\lambda_\nu^{\mathcal{B}}$ correspond to a substantially more negative impact. The results in panel (b) for $\lambda_\xi^{\mathcal{L}}$ and $\lambda_\xi^{\mathcal{B}}$, which capture the valuation for premium seating, are similar.

Taken together, the model provides a number of insights on how increasing consumer awareness of upgrade opportunities can impact profitability. In particular, we demonstrate the possibility that such efforts can reduce profits. However, the effect can differ substantially depending on the composition of consumers in a market.

3 Data

In this section, we describe our data sources and present a descriptive overview of participation in the upgrade processes during the sample and the impact of increased consumer awareness on seat allocations and profits.

3.1 Data Sources & Summary Statistics

The data come from a major North American airline and include flight-level seat inventories, passenger-level bookings, and information on upgrade activity (i.e., bidding in the auction and check-in sales). The sample spans nine months prior to the COVID-19 pandemic, and the revenue-management system and check-in upgrades were established prior to this period and unchanged throughout. The upgrade auction was introduced at the beginning of the sample, while the email-notification system to increase consumer awareness was introduced in the sixth month.

The flight-level inventory data track seat availability in the economy and premium cabins. Specifically, we observe the remaining capacity in both cabins at a daily frequency leading up to departure. The passenger-level revenue data provide information on all ticket transactions, including each passenger's identification number, itinerary, original cabin class, and fare paid. The fare paid for a passenger's itinerary is broken down by the airline to the flight-segment level. This simplifies our analysis of connecting and round-trip itineraries because we directly observe a price for each flight.

Figure 5: Example Slider

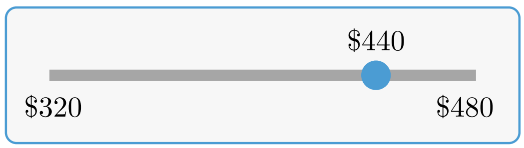

Notes: The figure shows an example of a slider seen by a passenger participating in an upgrade auction. The slider features a minimum, maximum, and starting position. Bids can be made by finalizing the position and submitting the bid among the discrete set of permissible values along the slider.

We observe detailed information on both upgrade processes. The first process allows economy passengers the option to purchase an upgrade to the premium cabin during check-in. Conditional on premium-cabin seat availability, these upgrades are allocated sequentially based on arrival for check-in. For upgrades purchased in this manner, we observe the passenger's identification number, itinerary details, and the fee paid.

The second process allows economy passengers who book at least one week in advance of departure to submit a bid to be upgraded to the premium cabin. Bids can be placed at any time between booking and the auction deadline using a slider interface, as illustrated in Figure 5. Passengers observe the initial slider position and finalize their bid by adjusting it among the discrete set of permissible values between the predefined minimum and maximum. For all submitted bids, we observe the passenger's identification number, itinerary details, bid amount, initial slider position, slider bounds, bid-acceptance decision, and the channel through which the bid was submitted (e.g., airline's app, web browser, etc).

The inventory, revenue, and upgrade data all include flight identifiers (i.e., departure date, directional segment, and flight number), which allows us to link the data sources seamlessly. This enables a complete description of the path taken by each passenger to a particular cabin on a flight (i.e., fare paid, initial cabin assignment, whether an upgrade was awarded, and if so through which process and at what cost).

In our analysis, we focus on flight numbers that were operated at least once per week and received at least one bid in the upgrade auction over the nine months.[9] This sample includes 62,481 flights from 279 unique flight numbers. Table 1 reports summary statistics in panel (a) at the flight level and in panel (b) at the passenger-transaction level.

[9]This removes flights operated on an irregular or seasonal basis, approximately 35% of all flights.

Table 1: Flight And Consumer Transaction Summary Statistics

(a) Flights

Variable	Mean	St. Dev.	25th	50th	75th	N
Flight Details						
Capacity Ratio (Econ/Prem)	10.9	2.2	10.2	10.2	13.5	62,481
Distance (Miles)	1019.1	664.2	427	800	1670	62,481
Load Factor at Departure						
Economy	0.815	0.187	0.735	0.883	0.951	62,481
Premium	0.817	0.247	0.667	0.917	1.000	62,481
Premium (Full Price)	0.581	0.297	0.333	0.583	0.833	62,481
Upgrades						
Auction Bids	1.3	3.2	0	0	1	62,481
Auction Upgrades	0.5	1.1	0	0	1	62,481
Check-In Upgrades	1.5	1.7	0	1	2	62,481
Upgrades	2.0	2.0	0	2	3	62,481

(b) Consumer Transactions

Variable	Mean	St. Dev.	25th	50th	75th	N
Bookings						
Economy	184.10	127.78	94.00	158.00	241.49	2,376,797
Premium	379.13	215.69	237.63	347.44	483.76	160,977
Upgrades						
Check-In Fee	104.29	80.70	39.00	79.00	149.00	93,975
Bid	113.33	103.53	35.00	70.00	160.00	46,542
Accepted Bid	96.60	85.67	35.00	60.00	140.00	31,710
Slider Minimum	92.92	87.88	25.00	50.00	140.00	46,542
Slider Maximum	188.31	141.36	80.00	125.00	260.00	46,542

Notes: Panel (a) and panel (b) present flight-level and transaction-level summary statistics, respectively.

Panel (a) shows that flights average 1,019 miles, with a substantial amount of variation across flights. The ratio of economy to premium seats is approximately eleven to one, but the two cabins have similar load factors of approximately 82% at departure. About 2 passengers are upgraded per flight, which is approximately 27% of those seated in the premium cabin at departure. This highlights the importance of these upgrade processes in allocating premium seating, which is an increasingly important contributor to airlines' profitability.[10]

Panel (b) shows that economy and premium fares average \$184.10 and \$379.13, respectively. The difference, which varies substantially across markets, demonstrates a strong preference among some passengers for premium seating. Check-in fees and accepted bids average \$104.29 and \$96.60, respectively, which are both less than the difference between the average economy and premium fares.[11] Although this is not necessarily problematic for the airline, e.g., accepted bids and check-in upgrades could occur primarily in markets with smaller differences in premium and economy fares, it highlights the possibility of a negative outcome for the airline from offering upgrade opportunities

[10] Regarding the importance of premium seating to airlines, see: www.reuters.com/airline-premium-seating.

[11] The average submitted bid is greater than the average accepted bid because larger bid values are more often placed in markets with a lower probability of acceptance.

to consumers.

As highlighted in Section 2, market heterogeneity has an important role in determining the effect of upgrade processes on profits. To create a classification of markets to guide our empirical analysis, we rely on the airline's internal taxonomy of markets that is similar to Li et al. [2014]. Specifically, we classify flights to Hawaii, Florida, Las Vegas, Arizona, California (excluding Los Angeles), and Mexico as *leisure* and flights to other destinations as *business*.[12]

Table 2: Flight And Consumer Transaction Summary Statistics: Market Heterogeneity

(a) Flights

	Leisure			Business		
Statistic	**Mean**	**Median**	**N**	**Mean**	**Median**	**N**
Flight Details						
Capacity Ratio (Econ/Prem)	12.3	13.5	9,382	10.7	10.2	53,099
Distance (Miles)	1559.5	1217	9,382	923.6	740	53,099
Load Factor at Departure						
Economy	0.891	0.938	9,382	0.801	0.869	53,099
Premium	0.861	1.000	9,382	0.809	0.917	53,099
Premium (Full Price)	0.735	0.833	9,382	0.554	0.583	53,099
Upgrades						
Auction Bids	1.6	0	9,382	1.2	0	53,099
Auction Upgrades	0.3	0	9,382	0.5	0	53,099
Check-In Upgrades	0.9	0	9,382	1.6	1	53,099
Upgrades	1.2	1	9,382	2.1	2	53,099

(b) Consumer Transactions

	Leisure			Business		
Statistic	**Mean**	**Median**	**N**	**Mean**	**Median**	**N**
Bookings						
Economy	211.57	188.48	344,577	179.44	151.41	2,032,220
Premium	468.84	427.08	26,996	361.05	329.00	133,981
Upgrades						
Check-In Fee	230.78	199.00	8,811	91.20	59.00	85,164
Bid	254.24	235.00	7,049	88.18	50.00	39,493
Accepted Bid	236.38	205.00	2,534	84.46	50.00	29,176
Slider Minimum	218.62	200.00	7,049	70.49	45.00	39,493
Slider Maximum	362.82	300.00	7,049	157.16	115.00	39,493

Notes: Panel (a) and panel (b) present flight-level and transaction-level summary statistics for each market type, respectively.

Table 2 presents the summary statistics from Table 1 by market type. Flights in leisure markets average 69% greater distance, operate at greater load factors, and have a greater proportion of premium seats purchased at full price than in business markets. Prices are greater on average in leisure markets, but less on a per-mile basis. Compared to leisure markets, the airline awards nearly one additional upgrade per flight in business markets (2.1 versus 1.2). Additionally, accepted bids

[12]Our data-use agreement limits the extent of the airline's network that we can disclose to ensure anonymity.

in business markets are substantially smaller in magnitude ($88 versus $254).

Figure 6: Distribution of Purchases and Average Fares Approaching Departure

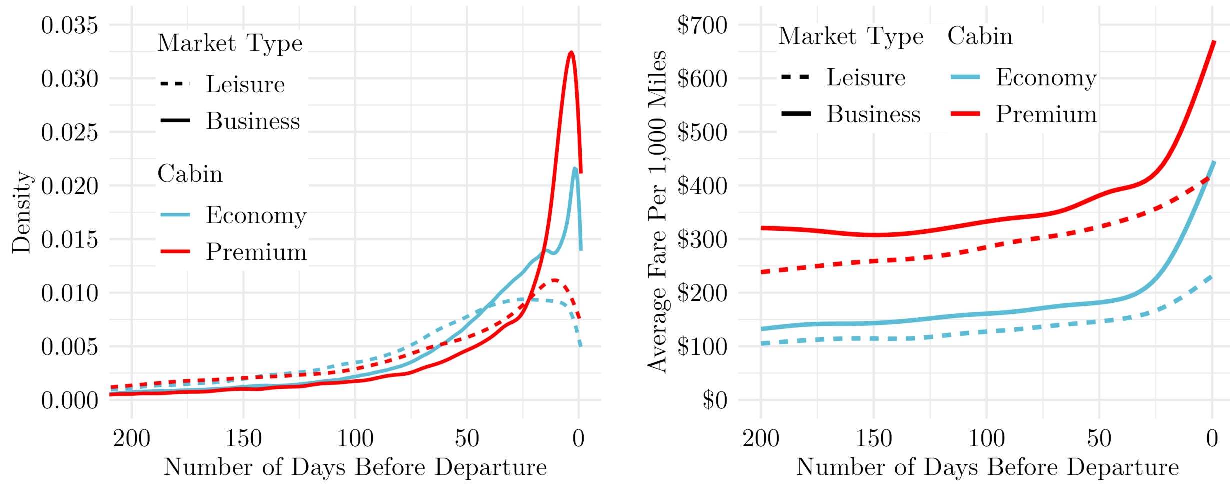

(a) Purchases Distribution

(b) Average Fares

Notes: Panel (a) presents the density of the timing of transactions relative to the departure date in business and leisure markets for both the premium and economy cabins. Panel (b) presents the average fare per thousand miles relative to departure date in business and leisure markets for both the premium and economy cabins.

To provide additional insight into the differences between market types, Figure 6 presents information on the timing of purchases and average fare relative to the departure date. Panel (a) presents the density of purchase timing for the economy and premium cabins by market type. Purchases in business markets are on average closer to departure than in leisure markets, 51.2 and 74.6 days, respectively. Premium purchases tend to occur closer to departure than economy in both types of markets. Panel (b) presents the path of average fares per thousand miles leading up to departure for both cabins by market type. Distance-adjusted fares in business markets are greater and rise more in the last 30 days before departure than in leisure markets. This pattern is consistent with inter-temporal price discrimination by the airline, resulting in substantially greater fares for late-arriving customers in business markets.

3.2 Participation & Capacity Allocation

There are several key differences between the upgrade auction and the check-in upgrade process. Check-in upgrades are made available for a fixed fee that is constant across flights on the same segment whenever there are unsold premium seats. The auction occurs a fixed number of days before departure, and passengers can submit bids at any time between booking and the date of the auction. In contrast to check-in upgrade offerings, bids can take on any value among the discrete values on the slider, which is constant across flights within a segment.

During the sample period, the airline introduced multiple channels through which passengers could access the slider and submit bids. Figure 7 illustrates participation in each of the channels during our sample period. Participation increased following the auction's implementation, until the

Figure 7: Auction Participation Over Time

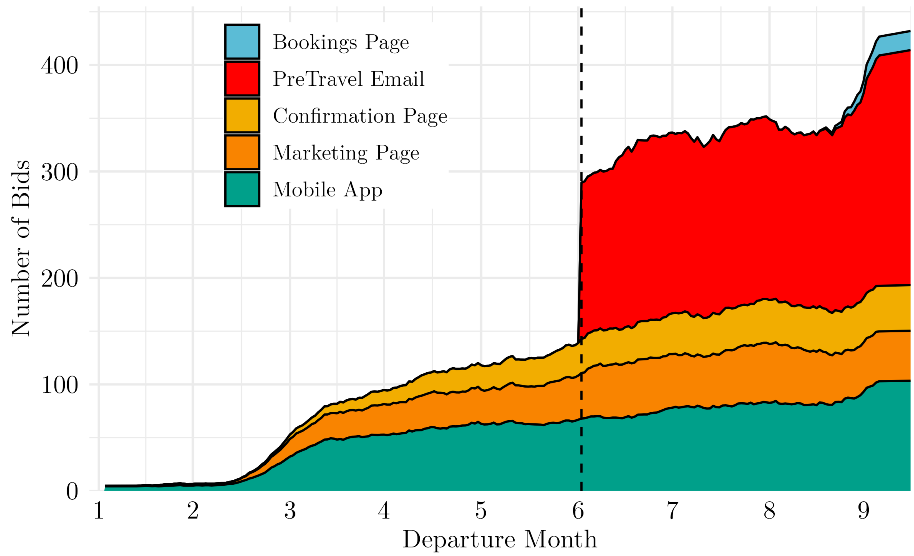

Notes: The figure shows the number bids by channel on each departure date. The airline adopted multiple channels during the sample, most notably the pre-travel email in the sixth month. We smooth the lines using a K-nearest neighbor kernel regression using $K = 7$ and a smoothing parameter of $h = 1.7$.

introduction of the pre-travel email notification, which immediately doubled the number of bids submitted per flight. The initial increase in participation at the beginning of the sample period slowed by the time the pre-travel email was introduced, suggesting that the subsequent increase in participation was driven primarily by improved passenger awareness.

The increase in participation provides the airline with additional opportunities to award consumers with upgrades and change the seating inventory. Panels (a) and (b) of Figure 8 show the average load factor leading up to departure in the economy and premium cabins, respectively, for all flights departing 60 days before and after the introduction of the pre-travel emails. By construction, a ticket purchase increases the load factor of the corresponding cabin, whereas an upgrade simultaneously decreases the economy load factor and increases the premium load factor, reflecting the reallocation of a passenger from economy to premium.

Before the introduction of the pre-travel email notifications, load factors increase in both cabins in the days leading up to check-in. At check-in, the economy load factor declines slightly, while the premium load factor increases sharply. After the introduction, these dynamics remain, but the effect of the auction becomes more evident. Specifically, there is a break in the upward trend of economy load factor and a concurrent jump in the premium load factor at the time of the auction. This pattern exists prior to the introduction, but is less pronounced due to fewer accepted bids.

The magnitude of the increase in load factor in the premium cabin, relative to the economy cabin, highlights a delicate trade-off for the airline. The airline must consider the possibility that each upgrade forecloses the possibility of an outright premium sale with no guarantee that the economy seat left by the upgraded passenger can be resold in the short time before departure with

Figure 8: Load Factors Approaching Departure

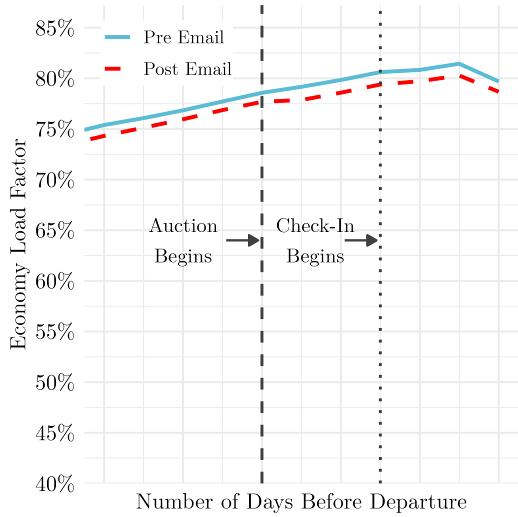

(a) Economy Cabin

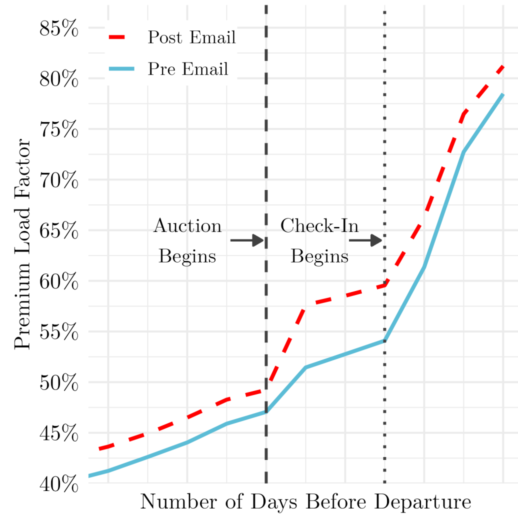

(b) Premium Cabin

Notes: Panel (a) presents the average load factor of the economy cabin relative to the departure date for flights departing 60 days before and 60 days after the email notifications began. Panel (b) presents the same calculation for the premium cabin.

waning demand for the economy cabin.[13]

In Section 2, we demonstrate that greater awareness and participation in the upgrade processes can have important implications for profits. The direct effect is an increase in revenue from accepted bids and check-in sales. Indirect effects arise from strategic consumers altering their initial purchase decisions to exploit the siloed nature of the pricing and upgrade teams' decision-making. This can reduce outright premium purchases and limit the airline's ability to re-sell the economy seats of upgraded passengers, while also altering the path of fares offered to customers. These different effects depend on the demand process, which make the sign and magnitude of the overall effect on revenue unclear.

Table 3 presents the average revenue per available seat (RAS) by source for flights with a departure date within 60 days before and after the introduction of the email notifications. The left panel presents the averages for all markets, while the middle and right panels present the analogous averages for business and leisure markets. We present results for revenue per available seat mile (RASM) in Table A.1 of the Appendix.[14]

The top panel of Table 3 presents the change to economy-cabin RAS. This is calculated as the total of fares for passengers seated in the economy cabin at departure divided by the total economy-cabin capacity across all flights in both the before and after periods. Across all markets, we find that economy RAS decreases by $3.52, which is consistent with an inability to re-sell the

[13] See Marsh et al. [2024a] for details on the arrival and search process of consumers on airline's websites and purchasing patterns by cabin.

[14] Given the fixed-cost nature of the airline industry, we focus on revenue as a measure of profitability and omit the peanut costs discussed Section 2. Given that upgrades increase the marginal cost of serving a passenger, this would overstate (understate) any increase (decrease) in profitability resulting from the email notifications.

Table 3: Change in Average Revenue Per Seat Before and After Pre-Travel Email

Market Type	All			Business			Leisure		
	Before	After	Δ	Before	After	Δ	Before	After	Δ
Econ Fare	116.92	113.41	-3.52	117.82	111.92	-5.90	110.99	121.60	10.61
Bids	2.84	7.08	4.25	2.60	6.86	4.27	4.42	8.30	3.88
Check-In Fee	12.55	12.00	-0.55	11.81	11.82	0.01	17.46	13.01	-4.45
Upgrade Fare	32.49	35.89	3.41	33.94	38.90	4.96	22.90	19.35	-3.55
Prem Fare	176.72	176.17	-0.54	168.11	160.47	-7.64	233.50	262.59	29.09
Total Revenue	122.92	119.99	-2.93	123.07	117.51	-5.55	121.95	133.60	11.65

Notes: The table presents the average RAS from different revenue sources for flights with a departure date within 60 days before and after the introduction of the email notifications. The left panel presents the statistics for all markets, while the middle and right panels present the same statistics for business and leisure markets.

seats of passengers awarded an upgrade through the auction. Economy RAS increases in leisure markets and decreases in business markets. An increase in economy RAS can arise if the option value associated with upgrades increases overall demand for the flight.

The middle panel of Table 3 presents the change to premium-cabin RAS, which we decompose into four sources: bid revenue, check-in fees, economy fares from upgraded passengers, and outright premium fare purchases. In the before and after periods, we divide the total of each source of revenue from passengers seated in the premium cabin at departure by the premium capacity across all flights in the respective period. We find that RAS from accepted bids and upgraded passenger fares increases, while RAS from check-in fees and outright premium fares decreases. Across all markets, premium RAS increases by \$6.57 after the notifications. In business markets, there is a lesser increase of \$1.60 to premium RAS, mainly due to less outright premium purchases. In contrast, premium RAS increases by \$24.97 in leisure markets, which can occur due to upward adjustments in fares resulting from additional accepted bids that create further scarcity of premium seats.

The bottom panel of Table 3 presents the change to total RAS. This is calculated by totaling all sources of revenue across all flights and dividing by the total capacity of the aircraft across all flights (i.e., economy and premium cabins). Given the high ratio of economy to premium seats, we find that total RAS decreases by \$2.93 after the introduction of the notifications. Total RAS decreases by \$5.55 in business markets and increases by \$11.65 in leisure markets. This heterogeneity is consistent with the model's predictions that even slight changes to demand primitives can alter the effects of increased awareness regarding upgrade opportunities.

Overall, the descriptive analysis shows substantive changes to purchasing behavior and revenues in response to increased awareness of and participation in the upgrade processes, and that these changes vary by the type of market.

4 Empirical Analysis

In this section, we enrich and expand on the descriptive analysis in Section 3 with a series of fixed-effects regressions using event-study and instrumental-variable (IV) frameworks to identify the effect of increased awareness regarding upgrade processes on profits.

4.1 Auction Participation and Profitability

To clarify the quantitative and statistical significance of the email-notification system on participation in and revenues from the upgrade processes, we estimate flight-level fixed-effects regressions of the form:

$$y_{it} = \beta \mathbb{1}[t \geq 0] + \alpha_i + \delta_{\text{dow}(t)} + \varepsilon_{it}. \tag{5}$$

We denote the outcome of interest for flight number i departing on date t by y_{it} (e.g., number of bids on a flight), fixed effects for flight numbers and departure day of week by α_i and $\delta_{\text{dow}(t)}$, respectively, and the error term by ε_{it}. Time is indexed so that $t < 0$ is prior to the pre-travel email and $t \geq 0$ is after the pre-travel email. This implies that the indicator, $\mathbb{1}[t \geq 0]$, equals one for flights departing after the introduction of the notification system, such that β measures the effect on y_{it}.

The estimates of β in Equation 5 for measures of participation in the upgrade processes are reported in Table 4.[15] For each outcome, we present results for 30-, 60-, and 90-day event windows (e.g., the 30-day window restricts the sample to flights departing within thirty days, before and after, the introduction of the email notifications) to explore the sensitivity of the results to the inherent trade-offs present in event studies.[16] Given the immediate and substantial impact of the notifications as shown in Figure 7, we expect these windows to adequately capture the entirety of any impact on consumer behavior and airline profits, with narrower windows decreasing the influence of confounding events.

Columns 1-5 of Table 4 present the results for the number of bids per flight, number of accepted bids per flight, bid revenue per flight, check-in revenue per flight, and submitted bid values, respectively. Given the similarity of the results across the event windows, we focus our discussion on the 60-day event window in the second panel of Table 4. The results for the 30-day and 90-day event windows are reported in the first and third panels, respectively.[17]

Columns 1 and 2 of Table 4 show that the number of submitted and accepted bids per flight increases by 0.86 and 0.55, respectively. Both results are statistically significant and economically meaningfully, i.e., both are approximately double the pre-notification levels. The implications of the increase in participation for upgrade revenues are reported in Columns 3 and 4. Auction revenue increases by \$51.24, and check-in revenue decreases by \$6.85 per flight. Finally, Column 5

[15] All regressions are estimated using the `feols` function in the `fixest` R package [Bergé, 2018].
[16] See MacKinlay [1997] for a formal discussion of these trade-offs for economic and finance applications.
[17] The similarity of the results for different event windows is not surprising given that day-of-week and flight-number fixed effects control for any change to the composition of flights within the respective samples.

Table 4: Effect of Email Introduction on Auction Outcomes

	Submitted (1)	Accepted (2)	Bid Revenue (3)	Check-In (4)	Bid Value (5)
30 Day Window					
Post Email	0.791***	0.500***	44.116***	-3.911	-0.051
	(0.034)	(0.025)	(2.948)	(3.666)	(0.652)
Observations	14,769	14,769	14,769	14,769	9,588
R^2	0.201	0.177	0.161	0.208	0.913
60 Day Window					
Post Email	0.857***	0.550***	51.243***	-6.845*	1.320**
	(0.033)	(0.024)	(3.140)	(3.541)	(0.619)
Observations	29,334	29,334	29,334	29,334	18,968
R^2	0.194	0.175	0.152	0.194	0.913
90 Day Window					
Post Email	0.921***	0.607***	56.635***	-7.027**	1.164*
	(0.034)	(0.026)	(3.136)	(3.255)	(0.674)
Observations	43,810	43,810	43,810	43,810	27,938
R^2	0.195	0.180	0.146	0.188	0.911
Fixed-effects					
Flight Number	Yes	Yes	Yes	Yes	Yes
Day of Week	Yes	Yes	Yes	Yes	Yes

Clustered (Flight Number) S.E. in parentheses
*Significance Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: The table presents estimates from Equation 5 using OLS with flight-number and day-of-week fixed effects. The dependent variables are, respectively, the number of bids per flight, number of accepted bids per flight, auction revenue per flight, check-in revenue per flight, and submitted bid value. Standard errors are clustered at the flight-number level. The panels present estimates with the sample restricted to flights within 30-, 60-, and 90-day windows of the the introduction of the notifications.

shows that the average submitted bid changes by only \$1.32. This is consistent with our modeling assumption in Section 2, i.e., low participation is driven primarily by awareness, not selection based on one's valuations of premium travel. Thus, the increase in auction revenue is driven almost entirely by a greater number of bids, not an increase in the value of submitted bids.

To understand the revenue implications from increased awareness of and participation in the upgrade processes, we examine measures of cabin-specific and total per-flight revenue using the same regression framework. We focus our discussion on revenue per available seat (RAS), which is calculated as described in Section 3. Economy RAS equals the total of economy fares paid by passengers that were not upgraded divided by the number of economy seats on the aircraft. Premium RAS equals the total of premium fares paid, economy fares of passengers that were upgraded, and upgrade revenue from accepted bids and check-in sales divided by the number of premium seats on the aircraft. Total RAS is simply the sum of all revenues divided by the total number of seats on the aircraft.

The estimates of β in Equation 5 for these measures of revenue are reported in Table 5. Like above, we report results for 30-, 60-, and 90-day event windows for the cabin-specific and total

Table 5: Effect of Email Introduction on Cabin Revenues Per Seat Flown

	Economy (1)	Premium (2)	Total (3)
30-Day Window			
Post Email	-2.74***	5.28**	-2.03*
	(1.02)	(2.31)	(1.04)
Observations	14,769	14,769	14,769
R^2	0.6038	0.5114	0.6253
60-Day Window			
Post Email	-3.66***	4.78*	-2.77**
	(1.03)	(2.44)	(1.07)
Observations	29,334	29,334	29,334
R^2	0.5762	0.4922	0.5999
90-Day Window			
Post Email	-2.76***	3.60	-1.95*
	(0.99)	(2.49)	(1.05)
Observations	43,810	43,810	43,810
R^2	0.5642	0.4804	0.5889
Fixed-effects			
Flight Number	Yes	Yes	Yes
Day of Week	Yes	Yes	Yes

Clustered (Flight Number) S.E. in parentheses
*Significance Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: The table presents estimates from Equation 5 using OLS with flight-number and day-of-week fixed effects. The dependent variables are RAS in the economy cabin, premium cabin, and overall. Standard errors are clustered at the flight-number level. The panels present estimates with the sample restricted to flights within 30-, 60-, and 90-day windows of the the introduction of the notifications.

revenue measures. Focusing on the 60-day window, we find that economy RAS decreased by \$3.66 and premium RAS increased by \$4.78. Total RAS decreased by \$2.77 due to the relatively greater number of economy seats. These results are all statistically significant, although the premium increase only at the 10% level. For a common Boeing 737 configuration of 12 premium and 132 economy seats, this corresponds to a decrease in total revenue per flight of \$425.76. The per-flight decrease in revenue using the results from the 30-day and 90-day event windows are qualitatively similar but about 25% smaller in absolute magnitude (i.e., \$298.32 and \$321.12, respectively).

In Table A.2 of the Appendix, we report results using revenue per available seat mile (RASM) as the dependent variable. The numerator of this measure is identical to RAS, but the denominator is the product of miles flown and available seats for each flight. Given the inclusion of flight number fixed effects in our analysis that control for any flight-specific heterogeneity, such as distance, we find similar results.

The results in Table 5 provide an indirect estimate of the relationship between participation in the upgrade processes and revenue, via the introduction of the notification emails. To provide a direct estimate of the relationship between the number of bids and revenues, which can be useful

Table 6: Effect of Number of Bids on Cabin Revenues Per Seat Flown

	Economy (1)	Premium (2)	Total (3)
30-Day Window			
Number of Bids	-3.457***	6.669**	-2.564*
	(1.314)	(2.915)	(1.330)
Observations	14,769	14,769	14,769
R^2	0.604	0.511	0.625
F-stat (1st stage)	1,303.345	1,303.345	1,303.345
60-Day Window			
Number of Bids	-4.274***	5.578**	-3.232**
	(1.231)	(2.820)	(1.272)
Observations	29,334	29,334	29,334
R^2	0.576	0.492	0.600
F-stat (1st stage)	2,946.626	2,946.626	2,946.626
90-Day Window			
Number of Bids	-3.001***	3.914	-2.120*
	(1.087)	(2.685)	(1.146)
Observations	43,810	43,810	43,810
R^2	0.564	0.480	0.589
F-stat (1st stage)	5,139.550	5,139.550	5,139.550
Fixed-effects			
Flight Number	Yes	Yes	Yes
Day of Week	Yes	Yes	Yes

Clustered (Flight Number) standard-errors in parentheses
*Significance Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: The table presents estimates from Equations 6-7 using two-stage least squares with the pre-travel email as an instrument for the number of bids. The dependent variables are RAS in the economy cabin, premium cabin, and overall. Standard errors are clustered at the flight-number level. The panels present estimates with the sample restricted to flights within 30-, 60-, and 90-day windows of the the introduction of the notifications.

for understanding the implications of efforts that increase participation by a given amount, we estimate the following system of equations

$$y_{it} = \alpha_i + \beta N_{it}^b + \delta_{\text{dow}(t)} + \varepsilon_{it}, \tag{6}$$
$$N_{it}^b = \omega_i + \theta \mathbb{1}[t \geq 0] + \gamma_{\text{dow}(t)} + \nu_{it}, \tag{7}$$

via two-stage least squares. Like above, y_{it}, α_i/ω_i, and $\delta_{\text{dow}(t)}/\gamma_{\text{dow}(t)}$ denote the per-flight outcome of interest (e.g., RAS) and flight-number and day-of-week fixed effects, respectively. In this framework, the introduction of the notifications is the instrument for the number of bids, N_{it}^b. Given that the email is not targeted or timed in a selective manner, it satisfies the necessary assumptions to influence auction participation, but not the underlying determinants of revenue (i.e., a valid instrument or IV).

Columns 1, 2, and 3 of Table 6 present the estimates of β from the IV regressions for economy, premium, and total RAS, calculated in the same manner as above. Focusing on the 60-day event window, we find that one additional bid results in a statistically significant decrease in economy and total RAS of $4.27 and $3.23, respectively, and an increase in premium RAS of $5.58. This corresponds to a loss of $496.68 for each additional bid on a flight using a common Boeing 737 configuration with 144 seats. The results are qualitatively similar but smaller in magnitude for 30-day ($376.30) and 90-day ($349.16) event windows. The analogous regression results using RASM as the dependent variable are presented in Table A.3 of the Appendix.

The estimates in Tables 5 and 6 provide consistent evidence that increased participation in the upgrade processes decreases total revenue on average. This is congruous with the predictions of the model in Section 2. Specifically, increases to efficiency and revenue from upgrades can be offset by the actions of strategic consumers when pricing and upgrade decisions are implemented independently, as in our data.

4.2 Market Heterogeneity

The model simulations in Section 2 suggest that our finding that revenue decreases on average with greater participation in the upgrade processes masks important heterogeneity across markets. In particular, given the parameters considered, the model predicts that increased awareness and participation in the upgrade processes in markets with a greater share of high-valuation travelers will have a more negative effect on profitability.[18]

Table 7 reports the estimates of Equation 5 by market type with RAS as the dependent variable.[19] In business markets, the notifications led to a decrease of $6.06 for economy RAS and increase of $1.67 for premium RAS, resulting in a statistically significant decrease of $5.13 in total RAS. This is consistent with the cannibalization of premium sales observed across all markets. In leisure markets, the notifications led to a statistically significant increase in both cabins, result in a statistically significant increase of $11.92 for total RAS. This is consistent with the option value associated with the auctions increasing demand for the flight, which can occur if customers that highly value premium service switch from the no-purchase option, another airline, or travel to another destination.[20]

Table 8 reports the IV estimates of Equation 6 by market type with RAS as the dependent variable.[21] In business markets, the 60-day window estimates imply each additional bid decreases economy RAS by $7.44 and increases premium RAS by $2.05 per seat, resulting in a statistically significant total decrease of $6.29. In leisure markets, an additional bid increases economy and premium RAS by $10.13 and $21.62, respectively, resulting in a statistically significant total increase of $10.72.

[18]Table A.4 in the Appendix reports the effect of the pre-travel email on auction participation by market type. In both cases, participation increased significantly, but the magnitude was larger in leisure markets.

[19]The RASM version of the regression results are in Table A.5 of the Appendix.

[20]In the conclusion, we discuss ways in which future research could extend the model to include competition and enrich the model of consumer choice.

[21]The RASM version of the regression results are in Table A.6 of the Appendix.

Table 7: Effect of Pre-Travel Email on Cabin Revenues Per Seat Flown by Market Type

Market Type	Business			Leisure		
Cabin	Economy (1)	Premium (2)	Total (3)	Economy (4)	Premium (5)	Total (6)
30-Day Window						
Post Email	-4.45***	2.63	-3.71***	8.67**	22.66***	9.17**
	(1.03)	(2.43)	(1.04)	(3.46)	(6.61)	(3.54)
Observations	12,805	12,805	12,805	1,964	1,964	1,964
R^2	0.63	0.52	0.65	0.32	0.31	0.35
60-Day Window						
Post Email	-6.06***	1.67	-5.13***	11.27***	24.05***	11.92***
	(1.00)	(2.56)	(1.05)	(3.47)	(6.95)	(3.56)
Observations	25,146	25,146	25,146	4,188	4,188	4,188
R^2	0.611	0.506	0.632	0.323	0.267	0.345
90-Day Window						
Post Email	-4.70***	1.02	-3.87***	8.96***	19.21***	9.66***
	(1.03)	(2.63)	(1.10)	(2.71)	(7.09)	(2.84)
Observations	37,340	37,340	37,340	6,470	6,470	6,470
R^2	0.60	0.49	0.62	0.28	0.25	0.30
Fixed-effects						
Flight Number	Yes	Yes	Yes	Yes	Yes	Yes
Day of Week	Yes	Yes	Yes	Yes	Yes	Yes

Clustered (Flight Number) standard-errors in parentheses
*Significance Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: The table presents estimates from Equation 5 using OLS with flight-number and day-of-week fixed effects. The dependent variables are RAS in the economy cabin, premium cabin, and overall. The left and right panels present estimates when the sample is restricted to business and leisure markets, respectively. Standard errors are clustered at the flight-number level. The panels present estimates with the sample restricted to flights within 30-, 60-, and 90-day windows of the the introduction of the notifications.

Taken together, these results provide strong empirical support for the model's predictions regarding heterogeneity in the profitability of the upgrade processes as implemented. The sign and magnitude of the effect depend critically on market-specific demand primitives. Business markets with greater heterogeneity in valuations and more dynamic changes in demand composition as departure approaches experience revenue decreases, while markets with more homogeneous and stable demand see revenue increases.

5 Conclusion

Revenue-management practices that dynamically adjust prices are well-established in many industries, but demand uncertainty can still lead to inefficient allocations and lost profit opportunities for firms (e.g., Aryal et al. [2023]). In settings with vertically-differentiated goods, such as travel and leisure markets, upgrade processes like auctions and fixed-price sales are widely used by firms seeking to increase efficiency and profitability.

Table 8: Effect of Number of Bids on Cabin Revenues Per Seat Flown by Market Type

Market Type	Business			Leisure		
Cabin	Economy	Premium	Total	Economy	Premium	Total
	(1)	(2)	(3)	(4)	(5)	(6)
30-Day Window						
Number of Bids	-5.78***	3.41	-4.83***	9.31**	24.31***	9.84**
	(1.36)	(3.16)	(1.38)	(3.64)	(6.67)	(3.70)
Observations	12,805	12,805	12,805	1,964	1,964	1,964
R^2	0.631	0.518	0.650	0.318	0.313	0.352
F-stat (1st stage)	1,138.777	1,138.777	1,138.777	171.751	171.751	171.751
60-Day Window						
Number of Bids	-7.44***	2.05	-6.29***	10.13***	21.62***	10.72***
	(1.27)	(3.13)	(1.32)	(3.17)	(6.09)	(3.24)
Observations	25,146	25,146	25,146	4,188	4,188	4,188
R^2	0.611	0.506	0.632	0.323	0.267	0.345
F-stat (1st stage)	2,490.532	2,490.532	2,490.532	476.156	476.156	476.156
90-Day Window						
Number of Bids	-5.25***	1.14	-4.32***	8.34***	17.87***	8.98***
	(1.17)	(2.93)	(1.24)	(2.67)	(6.41)	(2.77)
Observations	37,340	37,340	37,340	6,470	6,470	6,470
R^2	0.601	0.494	0.624	0.279	0.251	0.304
F-stat (1st stage)	4,442.311	4,442.311	4,442.311	738.074	738.074	738.074
Fixed-effects						
Flight Number	Yes	Yes	Yes	Yes	Yes	Yes
Day of Week	Yes	Yes	Yes	Yes	Yes	Yes

Clustered (Flight Number) standard-errors in parentheses
*Significance Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: The table presents estimates from Equations 6-7 using two-stage least squares with the pre-travel email as an instrument for the number of bids. The dependent variables are RAS in the economy cabin, premium cabin, and overall. The left and right panels present estimates when the sample is restricted to business and leisure markets, respectively. Standard errors are clustered at the flight-number level. The panels present estimates with the sample restricted to flights within 30-, 60-, and 90-day windows of the the introduction of the notifications.

We develop a model to demonstrate that strategic consumers can undermine the intended complementarity between revenue management and the upgrade processes, and that the overall impact on profits depends on features of the demand process. We test these predictions using unique data from a major airline that includes the introduction of an upgrade auction and an experiment with email notifications to increase awareness and bidding. This allows for clean identification of the effect of increased participation on the airline's profitability.

We find that, on average, increased awareness and participation in the upgrade processes leads to an increase in premium-cabin revenues, but a decrease in economy-cabin and total revenues. Consistent with the model's predictions, business markets are impacted most negatively. These results highlight the challenges faced by firms attempting to apply a siloed and uniform approach across the entirety of their networks with strategic consumers. Specifically, the upgrade processes must internalize its impact on consumers' initial cabin choices, rather than simply optimizing reallocation between cabins during the auction and at check-in.

Our paper leaves several avenues for future research. While our model and reduced-form approach effectively illustrate the underlying economics and measure the impact of a common form of implementation, we cannot fully explore the returns to further integration or alternative allocation strategies (e.g., integrating bid information into pricing policies, a single auction at departure to award premium seating, etc).[22] Further, airline pricing continues to evolve rapidly with the introduction of multiple levels of service quality, complex mixed bundling of ancillary offerings (e.g., bags and boarding priority), and personalized pricing that leverages vast data resources and machine-learning methodologies.[23] In these increasingly complex settings, the value and design of re-allocative strategies require further study.

[22] Marsh et al. [2024b] represents a first step in this direction by structurally estimating a model similar to ours on a set of concentrated markets, but computational tractability limits the range of counterfactual exercises and external validity to more competitive markets is limited.

[23] See Babii et al. [2025] for a discussion of personalized pricing in the airline industry and the effectiveness of targeted discounts.

References

Mark Armstrong. Recent developments in the economics of price discrimination. In Richard Blundell, Whitney K. Newey, and Torsten Persson, editors, *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, pages 97–141, 2006.

Gaurab Aryal, Charles Murry, and Jonathan W Williams. Price Discrimination in International Airline Markets. *The Review of Economic Studies*, page rdad037, 03 2023. ISSN 0034-6527. doi: 10.1093/restud/rdad037. URL https://doi.org/10.1093/restud/rdad037.

Andrii Babii, Alex Marsh, Garrett Scott, and Jonathan Williams. Personalized pricing in the airline industry. *Working Paper*, 2025.

Cynthia Barnhart and Yosef Sheffi. A Network-Based Primal-Dual Heuristic for the Solution of Multicommodity Network Flow Problems. *Transportation Science*, 27(2):102–117, May 1993. doi: 10.1287/trsc.27.2.102. URL https://ideas.repec.org/a/inm/ortrsc/v27y1993i2p102-117.html.

Cynthia Barnhart, Peter Belobaba, and Amedeo R. Odoni. Application of operations research in the air transport industry. *Transportation Science*, 37(4):368–391, 2003.

Laurent Bergé. *Efficient estimation of maximum likelihood models with multiple fixed-effects: the R package FENmlm*, 2018. URL https://cran.r-project.org/web/packages/fixest/index.html.

Jose M. Betancourt, Ali Hortaçsu, Aniko Oery, and Kevin R. Williams. Dynamic price competition: Theory and evidence from airline markets. *NBER Working Paper*, 4 2022. URL https://www.nber.org/system/files/working_papers/w30347/w30347.pdf.

Sungjin Cho, Gong Lee, John Rust, and Mengkai Yu. Optimal dynamic hotel pricing. *Mimeo*, 2018.

Pascal Courty and Hao Li. Sequential screening. *Review of Economic Studies*, 67(4):697–717, 2000.

Yao Cui, Izak Duenyas, and Ozge Sahin. Pricing of conditional upgrades in the presence of strategic consumers. *Management Science*, 64(7):3208–3226, 2018a. doi: 10.1287/mnsc.2017.2783. URL https://doi.org/10.1287/mnsc.2017.2783.

Yao Cui, Izak Duenyas, and Ozge Sahin. Unbundling of ancillary service: How does price discrimination of main service matter? *Manufacturing & Service Operations Management*, 20(3):455–466, 2018b. doi: 10.1287/msom.2017.0646. URL https://doi.org/10.1287/msom.2017.0646.

James D Dana. Using yield management to shift demand when the peak time is unknown. *RAND Journal of Economics*, 30(3):456–474, 1999.

Francesc Dilmé and Daniel F. Garrett. A dynamic theory of random price discounts. *Working Paper*, 12 2022. URL https://drive.google.com/file/d/1xOJ5SmArTDs9Kf-U5JB2fuTUeSrVZJmd/view.

Francesc Dilmé and Fei Li. Revenue Management without Commitment: Dynamic Pricing and Periodic Flash Sales. *The Review of Economic Studies*, 86(5):1999–2034, 12 2018. ISSN 0034-6527. doi: 10.1093/restud/rdy073. URL https://doi.org/10.1093/restud/rdy073.

Lingxiu Dong, Panos Kouvelis, and Zhongjun Tian. Dynamic pricing and inventory control of substitute products. *Manufacturing & Service Operations Management*, 11(2):317–339, 2009. doi: 10.1287/msom.1080.0221. URL https://pubsonline.informs.org/doi/abs/10.1287/msom.1080.0221.

Jeffrey C Ely, Daniel F. Garrett, and Toomas Hinnosaar. Overbooking. *Journal of European Economic Association*, 15(6):1258–1301, 2017.

Diego Escobari. Dynamic pricing, advance sales and aggregate demand learning in airlines. *The Journal of Industrial Economics*, 60(4):697–724, 2012.

Domenico Favrizi. Welfare consequences of upgrades: Evidence from the airline industry. *Working Paper*, 2024. URL https://www.domenicofabrizi.com/research.html.

I. Gale and T. Holmes. Advance-purchase discounts and monopoly allocation of capacity. *American Economic Review*, 83(1):135–146, 1993.

Guillermo Gallego and Ming Hu. Dynamic pricing of perishable assets under competition. *Management Science*, 60(5):1241–1259, 2014. doi: 10.1287/mnsc.2013.1821. URL https://doi.org/10.1287/mnsc.2013.1821.

Guillermo Gallego and Robert Phillips. Revenue management of flexible products. *Manuf. Serv. Oper. Manag.*, 6:321–337, 2004. URL https://api.semanticscholar.org/CorpusID:41335313.

Matthew Gentry and Tong Li. Identification in auctions with selective entry. *Econometrica*, 82(1):315–44, 2014. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/24029177.

Matthew Gentry, Tong Li, and Jingfeng Lu. Auctions with selective entry. *Games and Economic Behavior*, 105:104–111, 2017. ISSN 0899-8256. doi: https://doi.org/10.1016/j.geb.2017.06.016. URL https://www.sciencedirect.com/science/article/pii/S0899825617301148.

Alex Gershkov and Benny Moldovanu. Dynamic revenue maximization with heterogeneous objects: A mechanism design approach. *American Economic Journal: Microeconomics*, 1(2):168–98, August 2009. doi: 10.1257/mic.1.2.168. URL https://www.aeaweb.org/articles?id=10.1257/mic.1.2.168.

Kathryn Graddy and George Hall. A dynamic model of price discrimination and inventory management at the fulton fish market. *Journal of Economic Behavior & Organization*, 80(1):6–19, 2011.

Ali Hortaçsu, Olivia R Natan, Hayden Parsley, Timothy Schwieg, and Kevin R Williams. Organizational Structure and Pricing: Evidence from a Large U.S. Airline*. *The Quarterly Journal of Economics*, page qjad051, 09 2023. ISSN 0033-5533. doi: 10.1093/qje/qjad051. URL https://doi.org/10.1093/qje/qjad051.

Kinshuk Jerath, Serguei Netessine, and Senthil K. Veeraraghavan. Revenue management with strategic customers: The value of probabilistic selling. *Management Science*, 56(2):430–448, 2010. doi: 10.1287/mnsc.1090.1092. Explores probabilistic selling, including opaque upgrades, to manage consumer demand.

John Lazarev. The welfare effects of intertemporal price discrimination: An empirical analysis of airline pricing in u.s. monopoly markets. *Working Paper*, 2013.

Dan Levin and James L. Smith. Equilibrium in auctions with entry. *The American Economic Review*, 84(3):585–599, 1994. ISSN 00028282. URL http://www.jstor.org/stable/2118069.

Fei Li, Garrett Scott, and Jonathan W. Williams. Risk aversion and strategic timing of purchase in dynamic pricing environments: Evidence from the airline industry. *Working Paper*, 2019.

Jun Li, Nelson Granados, and Serguei Netessine. Are consumers strategic? structural estimation from the air-travel industry. *Management Science*, 60(9):2114–2137, 2014.

Tong Li and Xiaoyong Zheng. Entry and competition effects in first-price auctions: Theory and evidence from procurement auctions. *The Review of Economic Studies*, 76(4):1397–1429, 10 2009. ISSN 0034-6527. doi: 10.1111/j.1467-937X.2009.00558.x. URL https://doi.org/10.1111/j.1467-937X.2009.00558.x.

A. Craig MacKinlay. Event studies in economics and finance. *Journal of Economic Literature*, 35 (1):13–39, 1997.

Constantinos Maglaras and Joern Meissner. Dynamic pricing strategies for multiproduct revenue management problems. *Manufacturing & Service Operations Management*, 8(2):136–148, 2006. doi: 10.1287/msom.1060.0105. URL https://doi.org/10.1287/msom.1060.0105.

Vadim Marmer, Artyom Shneyerov, and Pai Xu. What model for entry in first-price auctions? a nonparametric approach. *Journal of Econometrics*, 176(1):46–58, 2013. ISSN 0304-4076. doi: https://doi.org/10.1016/j.jeconom.2013.04.005. URL https://www.sciencedirect.com/science/article/pii/S0304407613000821.

Alex Marsh, Garret Scott, Drew Van Kuiken, and Jonathan Williams. What can web traffic reveal about air-travel demand? *Working Paper*, 2024a.

Alex Marsh, Garrett Scott, and Jonathan W. Williams. Revenue management with reallocation. *Working Paper*, 2024b.

Roy Radner. The organization of decentralized information processing. *Econometrica*, 61(5):1109–1146, 1993. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/2951495.

James W. Roberts and Andrew Sweeting. When should sellers use auctions? *American Economic Review*, 103(5):1830–61, August 2013. doi: 10.1257/aer.103.5.1830. URL https://www.aeaweb.org/articles?id=10.1257/aer.103.5.1830.

William F. Samuelson. Competitive bidding with entry costs. *Economics Letters*, 17(1):53–57, 1985. ISSN 0165-1765. doi: https://doi.org/10.1016/0165-1765(85)90126-0. URL https://www.sciencedirect.com/science/article/pii/0165176585901260.

Garrett Scott. Screening with refunds: Evidence from the airline industry. *Working Paper*, 2024.

Zuo-Jun Max Shen and Xuanming Su. Customer behavior modeling in revenue management and auctions: A review and new research opportunities. *Management Science*, 53(12):1813–1830, 2007. doi: 10.1287/mnsc.1070.0786. Reviews customer behavior models relevant to designing upgrade opportunities.

Lifei Sheng, Xuying Zhao, and Christopher Thomas Ryan. Selling bonus actions in video games. *Management Science*, 71(3):2544–2564, 2025. doi: 10.1287/mnsc.2022.02348. URL https://doi.org/10.1287/mnsc.2022.02348.

Nancy Stokey. Intertemporal price discrimination. *Quarterly Journal of Economics*, 93(3):355–371, 1979.

Andrew Sweeting. Dynamic pricing behavior in perishable goods markets: Evidence from secondary markets for major league baseball tickets. *Journal of Political Economy*, 120(6):1133–1172, 2010.

Kalyan Talluri and Garrett van Ryzin. An analysis of bid-price controls for network revenue management. *Management Science*, 44(11):1577–1593, 1998. ISSN 00251909, 15265501. URL http://www.jstor.org/stable/2634901.

Kalyan Talluri and Garrett van Ryzin. Revenue management under a general discrete choice model of consumer behavior. *Management Science*, 50(1):15–33, 2004. doi: 10.1287/mnsc.1030.0147. URL https://doi.org/10.1287/mnsc.1030.0147.

Gustavo Vulcano, Garrett van Ryzin, and Costis Maglaras. Optimal dynamic auctions for revenue management. *Management Science*, 48(11):1388–1407, 2002.

Caio Waisman. Selling mechanisms for perishable goods: An empirical analysis of an online resale market for event tickets. *Quantitative Marketing and Economics*, 19:127–178, June 2021.

Kevin R. Williams. The welfare effects of dynamic pricing: Evidence from airline markets. *Econometrica*, 90(2):831–858, 2022. doi: https://doi.org/10.3982/ECTA16180. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA16180.

A Online Appendix

Table A.1: Change in Average Revenue Per Seat Mile Before and After Pre-Travel Email

Market Type	All			Business			Leisure		
	Before	After	Δ	Before	After	Δ	Before	After	Δ
Econ Fare	0.1554	0.1488	-0.0066	0.1669	0.1601	-0.0068	0.0797	0.0867	0.0070
Bids	0.0027	0.0070	0.0043	0.0027	0.0073	0.0046	0.0030	0.0055	0.0025
Check-In Fee	0.0133	0.0130	-0.0003	0.0133	0.0136	0.0003	0.0128	0.0095	-0.0033
Upgrade Fare	0.0426	0.0485	0.0059	0.0464	0.0547	0.0083	0.0171	0.0145	-0.0026
Prem Fare	0.2209	0.2096	-0.0112	0.2297	0.2146	-0.0151	0.1627	0.1825	0.0198
Total Revenue	0.1620	0.1554	-0.0066	0.1734	0.1664	-0.0070	0.0873	0.0950	0.0077

Notes: The table presents the average RASM from different revenue sources for flights with a departure date within 60 days before and after the introduction of the email notifications. The left panel presents the statistics for all markets, while the middle and right panels present the same statistics for business and leisure markets.

Table A.2: Effect of Email Introduction on Cabin Revenues Per Seat Mile Flown

	Economy (1)	Premium (2)	Total (3)
30-Day Window			
Post Email	-0.004***	0.001	-0.004***
	(0.001)	(0.003)	(0.001)
Observations	15,060	15,060	15,060
R^2	0.688	0.420	0.687
60-Day Window			
Post Email	-0.007***	-0.0022	-0.006***
	(0.001)	(0.0032)	(0.001)
Observations	29,972	29,972	29,972
R^2	0.683	0.400	0.682
90-Day Window			
Post Email	-0.006***	-0.0041	-0.006***
	(0.001)	(0.0032)	(0.001)
Observations	44,805	44,805	44,805
R^2	0.682	0.395	0.682
Fixed-effects			
Flight number	Yes	Yes	Yes
Day of week	Yes	Yes	Yes

Clustered (Flight Number) S.E. in parentheses
*Significance codes: *** 0.01, ** 0.05, * 0.1*

Notes: The table presents estimates from Equation 5 using OLS with flight-number and day-of-week fixed effects. The dependent variables are RASM in the economy cabin, premium cabin, and overall. Standard errors are clustered at the flight-number level. The panels present estimates with the sample restricted to flights within 30-, 60-, and 90-day windows of the the introduction of the notifications.

Table A.3: Effect of Number of Bids on Cabin Revenues Per Seat Mile Flown

	Economy (1)	Premium (2)	Total (3)
30-Day Window			
Number of Bids	-0.003***	0.001	-0.003***
	(0.001)	(0.002)	(0.001)
Observations	15,060	15,060	15,060
R^2	0.666	0.421	0.670
F-stat (1st stage)	679.768	679.768	679.768
60-Day Window			
Number of Bids	-0.005***	-0.002	-0.005***
	(0.001)	(0.002)	(0.001)
Observations	29,972	29,972	29,972
R^2	0.636	0.395	0.642
F-stat (1st stage)	1,358.783	1,358.783	1,358.783
90-Day Window			
Number of Bids	-0.004***	-0.003	-0.004***
	(0.001)	(0.002)	(0.001)
Observations	44,805	44,805	44,805
R^2	0.648	0.386	0.652
F-stat (1st stage)	2,500.832	2,500.832	2,500.832
Fixed-effects			
Flight number	Yes	Yes	Yes
Day of week	Yes	Yes	Yes

Clustered (Flight Number) standard-errors in parentheses
*Significance codes: *** 0.01, ** 0.05, * 0.1*

Notes: The table presents estimates from Equations 6-7 using two-stage least squares with the pre-travel email as an instrument for the number of bids. The dependent variables are RASM in the economy cabin, premium cabin, and overall. Standard errors are clustered at the flight-number level. The panels present estimates with the sample restricted to flights within 30-, 60-, and 90-day windows of the the introduction of the notifications.

Table A.4: Effect of Pre-Travel Email on Auction Participation by Market Type

Market Type	Business	Leisure
	Number of Bids	Number of Bids
	(1)	(2)
30-Day Window		
Post Email	0.77***	0.93***
	(0.04)	(0.10)
Observations	12,805	1,964
R^2	0.206	0.183
60-Day Window		
Post Email	0.82***	1.11***
	(0.03)	(0.09)
Observations	25,146	4,188
R^2	0.198	0.178
90-Day Window		
Post Email	0.90***	1.07***
	(0.04)	(0.09)
Observations	37,340	6,470
R^2	0.202	0.170
Fixed-effects		
Flight Number	Yes	Yes
Day of Week	Yes	Yes

Clustered (Flight Number) standard-errors in parentheses
*Significance Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: The table presents estimates from Equation 7 using OLS with flight-number and day-of-week fixed effects. The left and right columns present estimates when the sample is restricted to business and leisure markets, respectively. The dependent variable is the number of bids per flight. Standard errors are clustered at the flight-number level. The panels present estimates with the sample restricted to flights within 30-, 60-, and 90-day windows of the the introduction of the notifications.

Table A.5: Effect of Pre-Travel Email on Cabin Revenues Per Seat Mile Flown by Market Type

Market Type	Business			Leisure		
Cabin	Economy (1)	Premium (2)	Total (3)	Economy (4)	Premium (5)	Total (6)
30-Day Window						
Post Email	-0.005***	-0.001	-0.005***	0.006**	0.016***	0.006**
	(0.001)	(0.003)	(0.001)	(0.002)	(0.005)	(0.002)
Observations	12,805	12,805	12,805	1,964	1,964	1,964
R^2	0.655	0.398	0.654	0.516	0.375	0.542
60-Day Window						
Post Email	-0.009***	-0.005	-0.008***	0.008***	0.016***	0.009***
	(0.001)	(0.004)	(0.002)	(0.002)	(0.005)	(0.002)
Observations	25,146	25,146	25,146	4,188	4,188	4,188
R^2	0.649	0.376	0.648	0.491	0.338	0.515
90-Day Window						
Post Email	-0.008***	-0.007*	-0.007***	0.007***	0.012**	0.007***
	(0.001)	(0.004)	(0.002)	(0.002)	(0.005)	(0.002)
Observations	37,340	37,340	37,340	6,470	6,470	6,470
R^2	0.646	0.369	0.646	0.479	0.335	0.506
Fixed-effects						
Flight Number	Yes	Yes	Yes	Yes	Yes	Yes
Day of Week	Yes	Yes	Yes	Yes	Yes	Yes

Clustered (Flight Number) standard-errors in parentheses
*Significance Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: The table presents estimates from Equation 5 using OLS with flight-number and day-of-week fixed effects. The dependent variables are RASM in the economy cabin, premium cabin, and overall. The left and right panels present estimates when the sample is restricted to business and leisure markets, respectively. Standard errors are clustered at the flight-number level. The panels present estimates with the sample restricted to flights within 30-, 60-, and 90-day windows of the the introduction of the notifications.

Table A.6: Effect of Number of Bids on Cabin Revenues Per Seat Mile Flown by Market Type

Market Type	Business			Leisure		
Cabin	Economy (1)	Premium (2)	Total (3)	Economy (4)	Premium (5)	Total (6)
30-Day Window						
Number of Bids	-0.007***	-0.001	-0.006***	0.006**	0.017***	0.007**
	(0.002)	(0.004)	(0.002)	(0.002)	(0.005)	(0.003)
Observations	12,805	12,805	12,805	1,964	1,964	1,964
R^2	0.655	0.398	0.654	0.516	0.375	0.542
F-stat (1st stage)	1,138.777	1,138.777	1,138.777	171.751	171.751	171.751
60-Day Window						
Number of Bids	-0.007***	-0.004	-0.007***	0.004***	0.008***	0.004***
	(0.001)	(0.003)	(0.001)	(0.001)	(0.002)	(0.001)
Observations	25,146	25,146	25,146	4,188	4,188	4,188
R^2	0.649	0.376	0.648	0.491	0.338	0.515
90-Day Window						
Number of Bids	-0.009***	-0.007*	-0.008***	0.006***	0.011**	0.007***
	(0.002)	(0.004)	(0.002)	(0.002)	(0.005)	(0.002)
Observations	37,340	37,340	37,340	6,470	6,470	6,470
R^2	0.646	0.369	0.646	0.479	0.335	0.506
F-stat (1st stage)	4,442.311	4,442.311	4,442.311	738.074	738.074	738.074
Fixed-effects						
Flight Number	Yes	Yes	Yes	Yes	Yes	Yes
Day of Week	Yes	Yes	Yes	Yes	Yes	Yes

Clustered (Flight Number) standard-errors in parentheses
*Significance Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: The table presents estimates from Equations 6-7 using two-stage least squares with the pre-travel email as an instrument for the number of bids. The dependent variables are RASM in the economy cabin, premium cabin, and overall. The left and right panels present estimates when the sample is restricted to business and leisure markets, respectively. Standard errors are clustered at the flight-number level. The panels present estimates with the sample restricted to flights within 30-, 60-, and 90-day windows of the the introduction of the notifications.