Allocating Upgrades: Challenges and Opportunities in the Airline Industry*

Alex Marsh Garrett Scott Jonathan W. Williams

Last Updated: November 4, 2024

Preliminary: Not Do Circulate

Click here for current version.

Abstract

We study the allocation of premium-cabin upgrades through auctions and fixed-price sales at check-in. Our data comes from a major airline and includes information on ticket sales, aircraft inventory, and upgrade purchases and bids, before and after the introduction of the upgrade mechanisms. We use a model to identify challenges and trade-offs associated with these practices and highlight factors that can impact their effectiveness. As implemented, we find that these practices lead to a minimal increase in revenue because upgrade opportunities largely cannibalize outright premium-cabin sales. We show that information dissemination and framing offer meaningful opportunities to improve performance, and discuss how to avoid cannibalization through improved integration with existing revenue-management systems.

Keywords: Airlines, Auctions, Upgrades,

JEL Codes: L11, L21, L93

*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 Matt Leisten, Ben Rosa, and Timothy Hubbard for detailed discussions.

1 Introduction

Firms use a variety of strategies to allocate differentiated products to consumers with heterogeneous preferences. Increasingly often, firms employ sophisticated hybrid approaches that combine different strategies to further their profit-maximization objective by segmenting consumers based on preferences. The benefit of such approaches is not always ex-ante obvious, because it depends on the complementarity of the strategies in the particular setting and details of the integration.

One setting where these hybrid approaches are commonly used is travel and leisure markets, where a fixed inventory of vertically-differentiated goods and services are sold to a random number of sequentially arriving customers. In the case of air travel, airlines have traditionally varied prices approaching departure and across aircraft cabins to discriminate inter-temporally and screen passengers on their preference for quality. More recently, airlines have complemented these revenue-management practices with auctions and bundling. For example, auctions and fixed-price sales are used to award seat upgrades and ancillary services like checked bags are offered in bundles. In this paper, we study the impact of the introduction of upgrade processes on consumer choice and profitability of the airline, and provide insight into challenges to and opportunities for improving integration of these strategies with existing revenue-management practices.

Our empirical analysis relies upon a novel 15-month panel of data from a large North American airline that uses dynamic pricing to sell seats in economy and premium cabins. At the beginning of the sample, the airline began offering upgrades during check-in at a fixed market-specific price. Later in the sample, the airline gave customers that bought a base economy fare the option to bid for an upgrade to the premium cabin with decisions on acceptance made at a fixed date before departure. 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. No changes were made to the revenue-management system during this time, and the upgrade systems operated independently. Our sample includes information on every ticket purchase, daily aircraft inventory for each flight leading up to departure, and information on upgrade check-in sales and bidding. Additionally, we know the timing of practices introduced by the airline to raise awareness of the upgrade opportunities to customers (e.g., sending emails prior to check-in).

To provide insights into the trade-offs associated with this particular hybrid approach and guide our empirical analysis, we introduce and discuss the modeling framework from Marsh et al. [2024]. The model includes an airline that uses dynamic pricing to sell seats in two vertically-differentiated cabins before a fixed departure date, and equilibrium behavior among strategic consumers that account for upgrade opportunities during purchasing decisions. The model offers predictions regarding the impact on profits and consumer behavior from the introduction of upgrade procedures that we can compare to the observed outcomes. For example, the option to purchase or bid for an upgrade can diminish the effectiveness of prices for screening customers between the economy and premium cabins, resulting in an ambiguous impact on profit. The model also has more nuanced predictions regarding consumer behavior, like the role that dynamic pricing has in driving selection into the upgrade processes that can help guide efforts to improve integration. Further, we're able to identify behaviors of consumers that are not captured by the model but can have important consequences for profitability and can help enrich future modeling efforts.

REFRESH WITH RESULTS: We begin by documenting the changes to consumer behavior from the introduction of the upgrade processes. We find that prior to the auction being introduced, i.e., first 6 months of our sample, on average XX% of the premium cabin was occupied and approximately XX% of those passengers purchased an upgrade at check-in. The premium cabin was more than 80% occupied at departure for XX% of flights, which motivated many consumers to purchase outright rather than risk not being able to purchase an upgrade at check-in. The total price paid by passengers that bought premium outright was approximately XX% greater than those that upgraded (i.e., economy price plus check-in upgrade fee). Upon the introduction of the auction, we find that the fraction of premium seats that are occupied increases from XX% to XX%. Of these passengers, XX% were awarded an upgrade through the auction, XX% bought an upgrade at check-in, and XX% bought premium outright. Thus, the auction largely served to displace upgrade purchases at check-in. However, the average accepted bid is approximately XX% greater than the check-in price, which is consistent with the competitive nature of the auction increasing revenue.

REFRESH WITH RESULTS: The patterns in consumers bidding and the airline's acceptance of bids are consistent with predictions from the model. Overall, the airline is more likely to accept a bid if fewer premium seats are occupied at the time of the auction. However, the probability of winning is not monotonically increasing in the value of a bid within a market. A greater bid value is more attractive to the airline, but those bids tend to be placed when the difference between premium and economy fares is largest due to a small number of premium seats remaining. This selection into bidding due to the fluctuations in the dynamically set prices drives the non-monotonicity in win probabilities.

REFRESH WITH RESULTS: Taken together, these changes in consumer behavior make the profit implications theoretically ambiguous for airlines. Selection into the upgrade process when premium prices are greatest works against any increase in premium-seat occupancy. The impact on revenue from economy seat sales is also unclear. The direct effect of upgrades is negative, i.e., a seat is now empty in the economy cabin, but the option value associated with the upgrade process could increase economy purchases. To identify the impact of the introduction of the upgrade auction, we run a series of within-market regressions to measure how load factors and revenue changed in each cabin and overall. We find that premium revenue and load factor increased by XX% and XX%, respectively, while total revenue increased by XX%. Thus, the upgrade auctions had a limited impact on profitability as implemented.

REFRESH WITH RESULTS: We also observe consumer behaviors that are not captured by the model of Marsh et al. [2024], but have important implications for effective implementation of the upgrade auction. During the sample, the airline added multiple channels to submit bids (e.g., airline's phone app) and notifications or nudges as the auction date approached (e.g., emails). We find that these efforts increased participation in the auctions and the revenue collected from accepted bids. Interestingly, we also see that consumers try to glean information from the presentation of the auction. In particular, random variation in the initial position of the bid slider is strongly positively correlated with the submitted bid. This is consistent with a belief by consumers that the initial position conveyed information regarding the competitiveness of the auction when it did not.

Collectively, there are a number of insights from our results that are useful to firms for integrating these types of upgrade procedures with revenue-management practices. First, the two systems cannot be siloed and operate independently. In particular, the dynamic-pricing policies of the airlines must be adjusted to account for the option that the upgrade processes offers to consumers in order to maintain the discriminatory intent of prices. Integration can also go further by allowing prices to adjust based on the collection of bids at each point in time before departure. This would permit the airline to be more selective through higher premium prices if the collection of bids was favorable. Similarly, rejected bids could be used to personalize check-in upgrade prices, similar to a multi-round auction. We provide simple back-of-the-envelope estimates of the potential value from these improvements to integration.

Related Literature. Our paper contributes to the expansive literature on price discrimination. This literature includes theoretical studies that characterize the distributional implications like Varian [1985] and Bergemann et al. [2015], and empirical studies that measure the implications in different settings like Leslie [2004], Crawford and Shum [2006], Hendel and Nevo [2006], McManus [2007], Mortimer [2007], Nair [2007], Aryal and Gabrielli [2019]. There are also a number of empirical studies that specifically focus on inter-temporal price discrimination in the airline industry like Escobari [2012], Lazarev [2013], and Williams [2022]. Our study contributes to this literature by measuring the return for profit and of a hybrid approach that combines auctions and dynamic pricing.

The studies of price-discrimination most closely related to our paper are Cui et al. [2018], Aryal et al. [2023], and Marsh et al. [2024]. Cui et al. [2018] characterize optimal pricing of upgrades in a static setting where consumers make strategic purchasing decisions based on the possibility of being upgraded. We extend the analysis a dynamic setting and measure the actual performance of these upgrade procedures for an airline while providing insights to integration with revenue-management systems. Aryal et al. [2023] quantify the inefficiencies that result from airlines using dynamic pricing to discriminate both inter-temporally and intra-temporally (i.e., between cabins) by characterizing an efficient frontier of outcomes as in Bergemann et al. [2015]. Our study provides evidence on the success of airlines in eliminating some of those allocative inefficiencies by combining dynamic pricing with upgrade opportunities. Our study complements the modeling and analysis in Marsh et al. [2024]. Our focus is measuring the impact from the introduction of upgrade processes and identifying opportunities to improving integration with revenue-management practices. Marsh et al. [2024] estimate a structural model to perform counterfactual calculations to measure the return to some of the opportunities that we highlight.

Our paper also contributes to the literature that studies the behavior of firms when making complex decisions with limited resources. This typically requires heuristic approaches, or breaking apart the larger optimization problem into a series of less-complicated ones (e.g., Radner [1993]). This is common for airlines, as heuristic approximations are the foundation for otherwise intractable network routing and pricing decisions (e.g., Barnhart and Sheffi [1993] and Barnhart et al. [2003]). Hortaçsu et al. [2023] demonstrate the meaningful loss for airlines in delegating tasks associated with setting prices to different units within the firm. Like Marsh et al. [2024], our study provides empirical insights into the value from efforts to improve integration of two particular tasks, revenue-management practices that dynamically set prices and efforts to sell upgrades through auctions.

More generally, we also contribute to the growing literature on dynamic 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]) contributions. Like our work that extends a baseline models to include equilibrium bidding for upgrades, others extensions 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 multi-product firms (e.g., Aryal et al. [2023], Maglaras and Meissner [2006], Dong et al. [2009], and Talluri and van Ryzin [2004]). Our research also contributes to the endogenous entry literature on auctions [e.g. Samuelson, 1985, Levin and Smith, 1994, Marmer et al., 2013, Roberts and Sweeting, 2013, Gentry and Li, 2014, Gentry et al., 2017], which arises due to customer selection arising from dynamically set prices. Regarding the airline industry, Talluri and van Ryzin [1998] examine auctions as a way to price aircraft seats and Ely et al. [2017] to resolve over-bookings. To our knowledge, Marsh et al. [2024] is the only paper to examine the interaction between dynamic pricing and auctions.

Vulcano et al. [2002] characterize optimal auctions for allocating multiple units of a homogeneous good. While our setting is similar, we allow for consumer arrival and demand to be non-stationary and have multiple qualities of seats.

Alex's rewriting and extension of the above text. Vulcano et al. [2002] characterize optimal auctions for allocating multiple units of a homogeneous good. Similarly, Gershkov and Moldovanu [2009] characterize the optimal mechanism for allocating multiple units of heterogeneous goods with known qualities that cannot be reallocated and show how it can be implemented with a menu of dynamic prices. Although our setting is similar, there are some major differences that extends the allocation problem and allows for more realistic consumer preferences. First, we allow consumer arrival and demand to be non-stationary which is needed to capture upward moving prices approaching departure, a known fact about prices in the airline industry. Second, we have multiple qualities of seats that can be reallocated using upgrades. This generalizes the mechanism the airline can use to maximize profits. Lastly, we allow consumer preferences for quality to be heterogeneous which captures differences in willingness to pay for the higher quality cabin. Although preferences are more realistic, this significantly increases the difficulty of the allocation problem because consumers have private information in multiple dimensions and the traditional mechanism design approach in Vulcano et al. [2002] and Gershkov and Moldovanu [2009] becomes infeasible.

Framing and behavioral literature

2 Data

Our data come from a major North American airline that serves domestic and trans-border markets. The fifteen months of data consist of seat inventories from every passenger flight and all associated information on consumer transactions (i.e., itineraries, fares, and upgrades). Altogether, the data summarize the airline's revenue management process.

2.1 Data Sources

The inventory data track seat availability leading up to departure for every flight departing in the sample period. For each flight, we observe the remaining capacity and number of bookings by aircraft cabin at a daily frequency. The accompanying booking data include information on every ticket transaction, revealing each passenger's identification number, itinerary, original aircraft-cabin class, and fare paid.

The data also detail two passenger upgrade mechanisms. First, the airline uses a standard fixed fee at passenger check in that varies by flight segment. The check-in data contain all upgrades through the check-in mechanism, revealing each passenger's identification number, itinerary details, and the upgrade fee paid. Additionally, the airline implements an auction held prior to passenger check in. The auction data contain information from all bids placed by passengers, including the passenger's identification number, itinerary details, the bid amount, the minimum and maximum bid the passenger could submit, and the bid acceptance decision.

Because all data sets feature the flight details (i.e., departure date, directional segment, and flight number) observations can be matched across the data sets by a uniquely identified flight. Similarly, passengers in the booking data can be linked to both upgrade data sets via the combination of the passenger identification number and itinerary details.

The data span a period of fifteen months prior to the COVID-19 pandemic. The period begins with an overhaul of the revenue management process and the introduction of the check-in upgrade mechanism. Six months into the sample, the airline implemented the auction upgrade mechanism, allowing us to observe a period before the auction in addition to the roll out of the auction. Figure 1 provides an overview of the data timeline.

2.2 Upgrade Processes

The check-in upgrade mechanism has become commonplace in the airline industry. During the check-in period, economy passengers receive the option to pay a fixed price to have their

Figure 1: Data Timeline

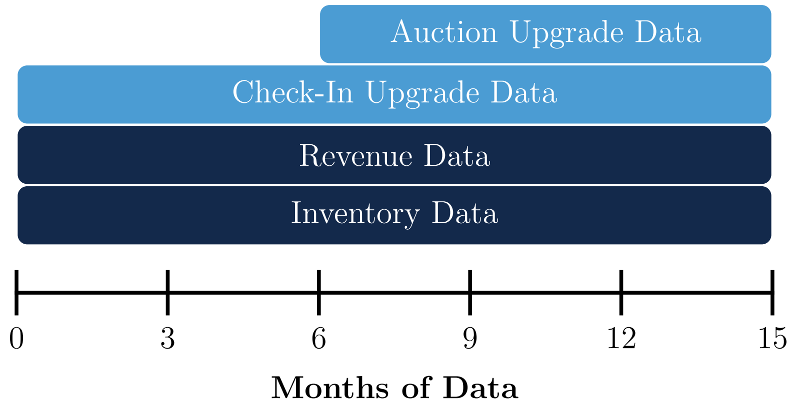

Notes: The figure details timeline of the data sources. The auction was introduced approximately six months into the sample.

ticket upgraded on their flight. The airline allocates upgrades sequentially based on check-in time, conditional on premium cabin seat availability.

The auction upgrade mechanism differs in both the allocation method and the timing of the upgrade. Economy passengers who book more than five days in advance of departure receive the option to place a bid to be upgraded on their flight. Bids can be placed any time between booking and five days before departure. With five days remaining, the airline selects which bids, if any, to accept.

Passengers place bids using a slider they are able to access on a web page, email, or phone app. Figure 2 displays a visual representation of the slider. The slider limits bids to a minimum and a maximum. Passengers observe the slider's initial position and place bids by finalizing the slider's position between the minimum and the maximum.

Figure 2: Example Slider

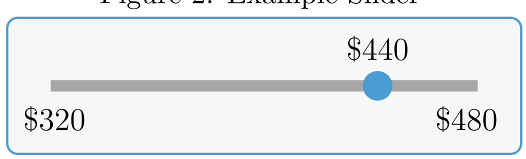

Notes: The figure shows an example of a slider seen by a passenger participating in an upgrade auction. The slider features a minimum, maximum, and starting position. Bids can be made by finalizing the position and submitting the bid.

The auction upgrade mechanism and the airline's revenue management system are not integrated. In fact, a third-party software manages the entire auction upgrade process. Furthermore, the revenue management system remained unchanged after the implementation of the auction. As a result, pricing decisions at the time of our sample did not reflect the change in the option value of the economy tickets associated with the new channel through which to move between cabins after booking.

2.3 Sample Selection & Descriptive Statistics

Our final sample includes all flights, bookings, and upgrades with a flight number that ever saw a bid placed for an upgrade auction. Table 1 shows basic descriptive statistics from the 177,108 flights in the sample. Of these flights, 74,755 were eligible for the auction. The mean and median flight distances in the sample are both approximately 1,200 miles, indicating a selection of medium-haul and long-haul flights. Across flights, the airline allocates capacity fairly consistently. Typically, there are between 10 and 14 economy seats for every premium seat, highlighting the opportunity cost associated with filling a seat in the premium cabin.

Table 1: Flight Statistics

Variable	**Mean**	**St. Dev.**	25^{th}	50^{th}	75^{th}	**N**
Flight Details						
Distance (Miles)	1220.060	760.518	451	1121	1742	177,108
Capacity (Econ./Prem.)	11.006	2.092	10.2	10.2	13.5	177,108
Load Factor (Departure)						
Economy	0.803	0.203	0.713	0.877	0.951	177,108
Premium	0.801	0.264	0.667	0.917	1.000	177,108
Premium (w/o Upgrades)	0.549	0.317	0.250	0.583	0.833	177,108
Upgrades						
Check-In Upgrades	2.030	2.201	0	1	3	177,108
Auction Bids	0.969	1.621	0	0	2	74,755
Auction Upgrades	0.630	0.476	0	0	1	74,755

Notes: The table shows descriptive statistics from flights in the sample. Details include the flight distances, seating capacities, allocation of seats, and allocation of upgrades.

Both cabins exhibit variation in the load factor (share of seating capacity sold) at departure. However, on average, approximately 20% of the seats in both the economy and premium cabins go unfilled. The variance in the load factor of the premium cabin at departure underscores the role upgrades serve in reducing the variability in occupancy. In fact, when isolating the share of seats allocated without upgrades, we find an inner-quartile range of 0.583. An average of 2.0 seats per flight are allocated in the premium cabin via upgrade at check in. Across the flights eligible for the auction, the airline receives almost one bid per flight and accepts approximately 65% of all submitted bids, resulting in an average of 0.6 seats allocated in the premium cabin via the upgrade auction.

Table 2 provides additional descriptive statistics from consumer transactions in the sample. The average price of an economy ticket is less than half the average premium ticket price, but there is substantial variation in the fare gap across markets, and within markets but across time.

Consumers pay $107.48 on average to be upgraded at check in. Alternatively, the average submitted bid in the upgrade auction is $136.09. In both cases, consumers pay less than the average price of an outright premium ticket, but those upgraded through the auction tend to pay more than those upgraded at check in. There is a clear trade-off associated with removing a premium seat from the inventory with time left to sell it outright. We find that the average economy cabin load factor at the time of the auction is 0.796. In contrast, the average premium cabin load factor at the same time is 0.506.

Table 2: Consumer Transaction Statistics

Variable	**Mean**	**St. Dev.**	25^{th}	50^{th}	75^{th}	**N**
Bookings						
Economy Price	182.40	127.60	96.50	153.98	234.00	18,038,507
Premium Price	398.09	218.92	251.00	367.00	509.00	1,100,438
Upgrades						
Check-In Fee	107.48	72.35	49.00	89.00	149.00	359,627
Submitted Bid	136.09	116.83	40.00	100.00	200.00	72,470
Slider Minimum	111.46	97.69	25.00	80.00	160.00	72,470
Slider Start	179.80	145.60	50.00	145.00	260.00	72,470
Slider Maximum	232.49	167.90	85.00	195.00	325.00	72,470

Notes: The table shows descriptive statistics from consumer transactions in the sample. These statistics come from bookings, as well as the check-in upgrade mechanism and the auction upgrade mechanism. Details include prices, fees, bids, and auction slider characteristics.

Given that upgrades account for more than a quarter of premium seat occupancy, understanding the trade-offs the airline and consumers face in the presence of upgrade mechanisms is crucial when considering the challenges and opportunities associated with using auctions alongside dynamic pricing. In the next section, we consider a model to quantify these trade-offs and help guide intuition.

3 Model

In this section, we use the modeling framework from Marsh et al. [2024] to identify the trade-offs of adding upgrade mechanisms to revenue management systems. The model features a monopoly airline seeking to maximize profits by allocating a fixed capacity of premium and economy seats on a flight. Utility-maximizing consumers arrive with individual preferences and choose from a set of options.

3.1 Timing

Time is discrete and denoted by $t \in \{1, ..., T, T+1\}$. The airline begins selling tickets for premium and economy seats in $t = 1$ and continues to sell tickets in each $t \leq T$. At the beginning of $\tilde{t}$, where $1 < \tilde{t} < T$, the airline upgrades consumers via the auction mechanism. In $T + 1$, consumers check in for their flight, and the airline upgrades consumers via the check-in mechanism. The flight departs at the end of time.

A flight has initial capacity $\boldsymbol{k}_1 = (k_1^f, k_1^e)$, which represent the total number of seats in the premium and economy cabins, respectively. Let $\boldsymbol{k}_t = (k_t^f, k_t^e)$ be the number of remaining seats at t. In $t \leq T$, the airline sets prices $\boldsymbol{p}_t = (p_t^f, p_t^e)$ for premium and economy tickets, respectively, before the arrival of any consumers.

3.2 Consumer Arrival and Preferences

N_t consumers arrive in period $t \leq T$. Upon arrival, consumers learn their preferences and choose an option from the choice set $\{f, e, o\}$, where f, e, and o represent the premium cabin, economy cabin, and outside option, respectively. Consumers cannot delay their decision. Consumer i's utility from choice m in period t is defined as

$$u_{it}^m \equiv \nu_i \xi_i^m - p_t^m, \tag{1}$$

where ν_i represents consumer i's willingness to pay for the flight and ξ_i^m is consumer i's quality measure of choice 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, allowing us to simply write $\xi_i^f = \xi_i$. Because the premium and economy cabins are vertically-differentiated, we allow the quality of premium seats to have support $\xi_i = [1, \infty)$.

Given prices and upgrade beliefs, consumer i chooses $m \in \{f, e, o\}$ that maximizes expected utility. In the absence of upgrade mechanisms, the consumer simply compares the utilities specified in Equation 1 and selects the available option that maximizes utility. However, in the presence of an upgrade mechanism, an economy ticket has an option value associated with the probability of being upgraded, while the utilities of choices f and o remain unchanged.

Before the introduction of the auction, consumer i's expected utility from choice e can be written as

$$U^e_{it} = u^e_{it} + \varphi_t \max\left\{0, \nu_i(\xi_i - 1) - r\right\}, \tag{2}$$

where u^e_{it} is the guaranteed utility from the economy cabin. Let φ_t be the belief about the probability a check-in upgrade is available at price r. Then, $\varphi_t \max\{0, \nu_i(\xi_i - 1) - r\}$ represents the expected utility gain from the check-in mechanism.

Alternatively, when both upgrade mechanisms are utilized, consumer i's expected utility from choice e is

$$\mathcal{U}^e_{it} = u^e_{it} + \varrho^*_{it}\big(\nu_i(\xi_i - 1) - b^*_{it}\big) + \varphi^a_t\big(1 - \varrho^*_{it}\big) \max\left\{0, \nu_i(\xi_i - 1) - r\right\}. \tag{3}$$

Again, u^e_{it} is the guaranteed utility from the economy cabin. Let ϱ^*_{it} be consumer i's belief about the probability of winning an upgrade through the auction given their optimal bid b^*_{it}. Then, $\varrho^*_{it}(\nu_i(\xi_i - 1) - b^*_{it})$ represents the expected utility gain from the auction mechanism. If consumer i does not win an upgrade through the auction, they believe a check-in upgrade with price r will be available with probability φ^a_t. Therefore, $\varphi^a_t(1 - \varrho^*_{it}) \max\{0, \nu_i(\xi_i - 1) - r\}$ represents the expected utility gain from the check-in mechanism.

3.3 Airline's Dynamic Program

In each $t \leq T$, the airline sets prices $\boldsymbol{p}_t = (p^f_t, p^e_t)$ prior to the arrival of the N_t consumers. Let $\mathbb{E}_t[Q^m(\boldsymbol{p}, \boldsymbol{k})]$ be the expected quantity demanded in cabin m at time t given prices $\boldsymbol{p}$ and remaining capacities $\boldsymbol{k}$. Then the airline's expected per-period revenues can be written $\mathbb{E}_t[R(\boldsymbol{p}, \boldsymbol{k})] = p^f \mathbb{E}_t[Q^f(\boldsymbol{p}, \boldsymbol{k})] + p^e \mathbb{E}_t[Q^e(\boldsymbol{p}, \boldsymbol{k})]$. Additionally, 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 $\boldsymbol{c} = (c^f, c^e)$. The airline realizes these costs at the end of $T + 1$. Total costs are defined as $C(\boldsymbol{k}) = c^f(k^f_1 - k^f_d) + c^e(k^e_1 - k^e_d)$, or the peanut costs from each consumer in the cabins at departure.

Because the airline sets the price of the check-in upgrade in advance of any sales, the airline's boundary condition is $V_{T+1}(\boldsymbol{k}) = 0$. For discount factor $\delta \in [0, 1]$, the value function in period T is

$$V_T(\boldsymbol{k}) = \max_{\boldsymbol{p} \in \mathbb{R}^2_+} \mathbb{E}_T[R(\boldsymbol{p}, \boldsymbol{k})] - \delta \int_{\boldsymbol{k}' \in \mathcal{K}} C(\boldsymbol{k}') dH_T(\boldsymbol{k}' | \boldsymbol{k}, \boldsymbol{p}), \tag{4}$$

where $\mathcal{K}$ is the state space and $H_t(\boldsymbol{k}' | \boldsymbol{k}, \boldsymbol{p})$ is the distribution of next period's state conditional upon this period's state and prices.

In period $t < T$, the airline's dynamic program is:

$$V_t(\boldsymbol{k}) = \max_{\boldsymbol{p}\in\mathbb{R}^2_+} \mathbb{E}_t[R(\boldsymbol{p},\boldsymbol{k})] + \delta \int_{\boldsymbol{k}'\in\mathcal{K}} V_{t+1}(\boldsymbol{k}')dH_t(\boldsymbol{k}'|\boldsymbol{k},\boldsymbol{p}). \tag{5}$$

The solution to Equations 4 and 5 is a policy function $\boldsymbol{p}_t(\boldsymbol{k})$, obtained using backwards induction.

3.4 Bid Acceptance

Let $\boldsymbol{i}^u = (-1, 1)$ be the upgrade vector (i.e. removing one seat from remaining premium capacity and adding one seat back to remaining economy capacity). Then, accepting n bids at $\tilde{t}$ implies a change in the capacity state from $\boldsymbol{k}_{\tilde{t}}$ to $\boldsymbol{k}_{\tilde{t}} + n\boldsymbol{i}^u$. The cost of accepting n bids at time t is the opportunity cost of selling those seats in the future with price or fixed fee mechanisms. As $V_t(\boldsymbol{k})$ is the value of having $\boldsymbol{k}$ seats at t and selling them only with dynamic pricing, let $U_t(\boldsymbol{k})$ be the value of having k^f premium seats at t and selling them through check-in upgrades. Because optimal prices are assumed to ignore upgrades, $V_t(\boldsymbol{k})$ and $U_t(\boldsymbol{k})$ are linearly separable and the total value of $\boldsymbol{k}$ in t can be written as $TV_t(\boldsymbol{k}) = V_t(\boldsymbol{k}) + U_t(\boldsymbol{k})$, and the total cost of n upgrades through the auction is the change in the total value function $TV_t(\boldsymbol{k}_t + n\boldsymbol{i}^{\boldsymbol{u}}) - TV_t(\boldsymbol{k}_t)$. Define the marginal opportunity cost of the n^{th} upgrade, denoted $\Delta TV_t(n, \boldsymbol{k})$, to be

$$\Delta TV_t(n,\boldsymbol{k}) = \begin{cases} 0 & \text{if } n = 0 \\ TV_t(\boldsymbol{k} + n\boldsymbol{i}^{\boldsymbol{u}}) - TV_t(\boldsymbol{k} + (n-1)\boldsymbol{i}^{\boldsymbol{u}}) & \text{if } n \in \{1, 2, ..., k^f\} \\ \infty & \text{otherwise.} \end{cases} \tag{6}$$

At the time of the auction in $\tilde{t}$, the airline orders all bids submitted in previous periods. The marginal revenue of the n^{th} upgrade is simply $b^{(n)}$. If $\Delta TV_t(n, \boldsymbol{k})$ is increasing in n, then the airline will continue accepting bids until the marginal cost of an upgrade exceeds the marginal revenue, or

$$\Delta TV_{\tilde{t}}(n+1, \boldsymbol{k}) > b_{\tilde{t}}^{(n+1)}. \tag{7}$$

Let n^u be the smallest n that satisfies Equation (7) given remaining capacities $\boldsymbol{k}$ and the set of bids (i.e. the number of bids accepted).

4 Empirical Analysis

The model in Section 3 introduces a set of trade-offs for consumers and the airline that arise with the implementation of the auction upgrade mechanism. In this section, we highlight some of the key trade-offs and use the data to explore the implications for consumer behavior and airline profits. We also consider consumer behavior not explicitly captured by the model that our data are uniquely able to address.

4.1 Consumer Response

First, the auction elicits a change in the option value to consumers purchasing economy tickets, as highlighted in Equations 2 and 3. Although both expected utilities are weakly greater than the utility from economy in the case of no upgrades,

Once the auction is implemented, an increase in sold capacity is observed in the premium cabin at five days out whereas the rate of increase in the economy cabin decreases. This comes from bids being accepted at five days out and economy passengers being moved to the premium cabin.

Figure 3: Share of Seats Sold Approaching Departure by Cabin

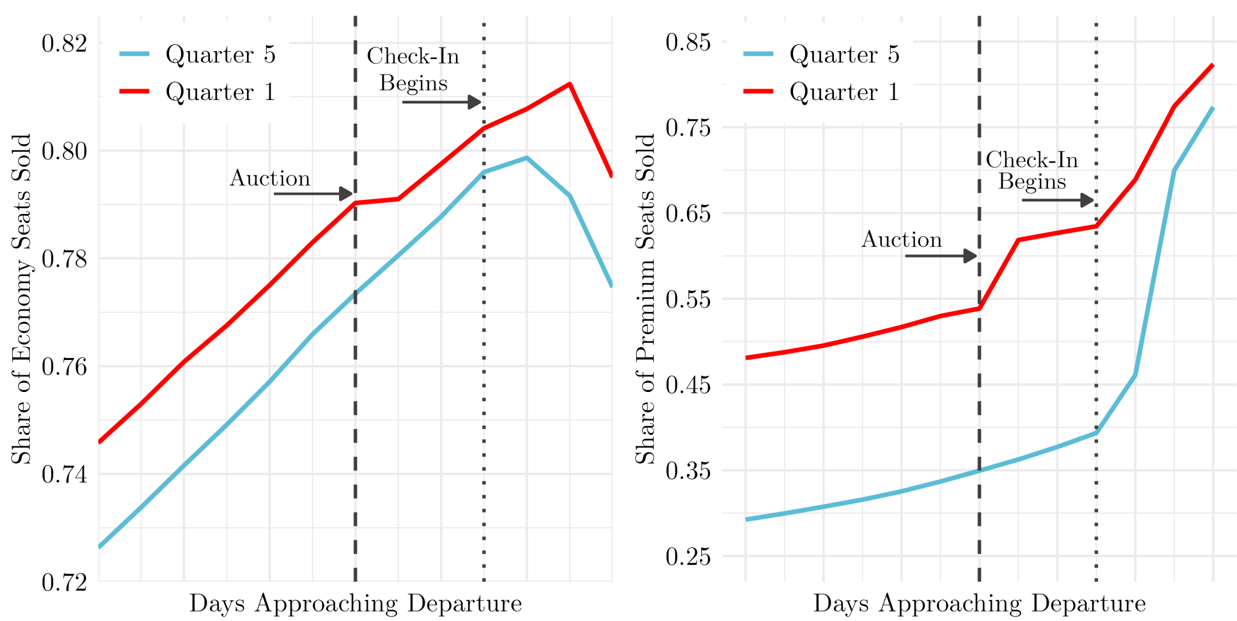

(a) Economy Cabin (b) Premium Cabin

Notes: Quarter 1 corresponds to months 1-3 and Quarter 5 corresponds to months 13-15.

Table 6 shows how bids are affected based on how close departure was when the bids were placed. While there is some evidence that

Table 3: Participation Regressions

	Submitted Bids Poisson Regression		
	(1)	(2)	(3)
Variables			
Load Factor at Auction: Economy	0.5874***		0.5793***
	(0.0356)		(0.0356)
Load Factor at Auction: Premium	0.2097***		0.2100***
	(0.0158)		(0.0158)
Slider Minimum		-0.0019***	-0.0018***
		(0.0004)	(0.0004)
Slider Maximum		-0.0013***	-0.0011***
		(0.0003)	(0.0003)
Fixed-effects			
Flight Number	Yes	Yes	Yes
Departure Date	Yes	Yes	Yes
Departure Day of Week	Yes	Yes	Yes
Fit statistics			
Observations	30,413	30,413	30,413
Pseudo R^2	0.07024	0.06359	0.07069
BIC	113,542.2	114,275.8	113,513.9

Clustered (Flight Number) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Notes: Estimated with the `fepois()` function in the `fixest` R package [Bergé, 2018].

Table 4: Participation Regressions

	Submitted Bids Poisson Regression			
	(1)	(2)	(3)	(4)
Variables				
Load Factor at Auction: Economy	0.5793***	0.3649***	0.6255***	0.1942***
	(0.0356)	(0.0514)	(0.0355)	(0.0502)
Load Factor at Auction: Premium	0.2100***	0.5826***	-0.0380*	0.5584***
	(0.0158)	(0.0276)	(0.0192)	(0.0274)
Slider Minimum	-0.0018***	-0.0017***	-0.0018***	-0.0018***
	(0.0004)	(0.0004)	(0.0004)	(0.0004)
Slider Maximum	-0.0011***	-0.0010**	-0.0009**	-0.0008*
	(0.0003)	(0.0003)	(0.0003)	(0.0003)
Departure Load Factor: Economy		0.3037***		0.6279***
		(0.0448)		(0.0462)
Departure Load Factor: Premium - Full		-0.4913***		-0.9046***
		(0.0311)		(0.0357)
Check-In Upgrades			-0.0700***	-0.0951***
			(0.0026)	(0.0029)
Fixed-effects				
Flight Number	Yes	Yes	Yes	Yes
Departure Date	Yes	Yes	Yes	Yes
Departure Day of Week	Yes	Yes	Yes	Yes
Fit statistics				
Observations	30,413	30,413	30,413	30,413
Pseudo R^2	0.07069	0.07302	0.07831	0.08546
BIC	113,513.9	113,277.4	112,683.6	111,914.7

Clustered (Flight Number) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Notes: Estimated with the `fepois()` function in the `fixest` R package [Bergé, 2018].

Table 5: Number of Check-Ins Purchased Regressions

	Poisson Regression	
	Check-In Upgrades	
	(1)	(2)
Variables		
Load Factor at Auction: Economy	0.5675***	-0.3795***
	(0.0325)	(0.0478)
Load Factor at Auction: Premium	-1.708***	-0.2929***
	(0.0227)	(0.0266)
Submitted Bids	-0.1013***	-0.1115***
	(0.0067)	(0.0068)
Auction Upgrades	-0.0314***	-0.0400***
	(0.0073)	(0.0074)
Departure Load Factor: Economy		1.406***
		(0.0452)
Departure Load Factor: Premium - Full		-1.808***
		(0.0331)
Fixed-effects		
Flight Number	Yes	Yes
Departure Date	Yes	Yes
Departure Day of Week	Yes	Yes
Fit statistics		
Observations	74,671	74,671
Pseudo R^2	0.19176	0.22553
BIC	273,625.3	262,702.4

Clustered (Flight Number) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Notes: Estimated with the `fepois()` function in the `fixest` R package [Bergé, 2018]. 84 observations were removed due to only 0 outcomes.

Table 6: Bidding Regressions: Approaching Departure

	Bid (1)	Normalized Bid (2)	Bid (3)	Normalized Bid (4)
Variables				
Slider Minimum	0.980***		0.972***	
	(0.027)		(0.027)	
Slider Maximum	0.043		0.034	
	(0.028)		(0.028)	
Number of Days Out	0.069***	0.057***	0.082***	0.065***
	(0.009)	(0.006)	(0.010)	(0.006)
Economy Fare Paid			0.037***	
			(0.004)	
log(Economy Fare Paid)				0.048***
				(0.004)
Fixed-effects				
Flight Number	Yes	Yes	Yes	Yes
Departure Date	Yes	Yes	Yes	Yes
Departure Day of Week	Yes	Yes	Yes	Yes
Fit statistics				
Observations	72,467	72,467	67,521	67,521
R^2	0.908	0.063	0.909	0.071

Clustered (Flight Number) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Notes: Estimated with the `feols()` function in the `fixest` R package [Bergé, 2018].

4.2 Implementation Impact

Figure 4: Bid Acceptance Probability by Remaining Capacity State

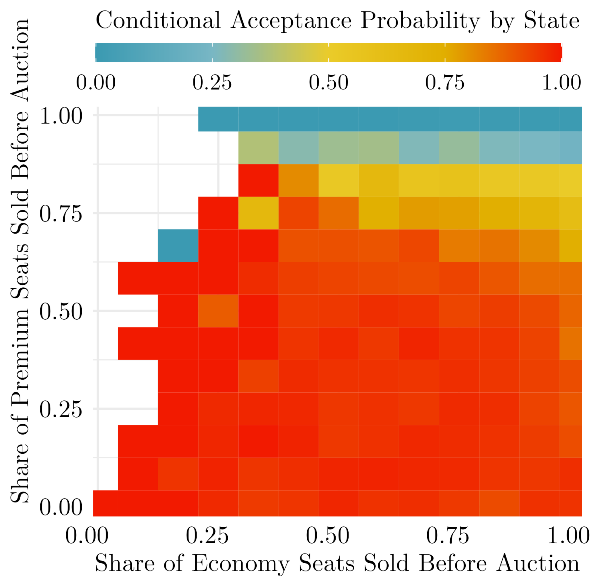

(a) Conditional Probability Heat-Map

(b) Joint Probability Heat-Map

Notes: A "state" is an ordered pair of the share of remaining seats (capacity) in each cabin the day before the auction. A tile on the heat-map represents one possible state. The economy and premium remaining capacity shares are binned by rounding to the nearest $1/12^{\text{th}}$. Figure 4b is the product of Figure 4a and the probability of each state (tile), and the probabilities in each tile sum up to the unconditional acceptance probability in Table 2. Figure 12 in the Appendix shows the distribution of states.

As fewer premium seats remain the day before the auction, the probability of being upgraded decreases. This is because the opportunity cost associated with an upgrade is larger the fewer premium seats remain. As the share of economy seats sold the day before the auction increases, a slight decreasing in the win probability can be observed for tiles in between 1 and 5 on the y-axis. This is likely due to two effects: a competition effect and demand effect. Flights that have sold more economy seats are going to have more bidders on average. This is the competition effect. However, flights that have sold more economy seats before the auction are likely to have higher overall demand, which means more consumers will arrive closer to departure willing to pay full price for premium, increasing the opportunity cost of an upgrade. This is the demand effect.

Economy and premium being full is the modal outcome and the distribution is concentrated in the bottom right corner.

Figure 5 shows the total premium capacity sold approaching departure broken out by how the seats were sold: at full price, with the auction, and at check-in. While a general increase in the number of premium seats sold outright (i.e. at full-price) from year 1 to year 2 is present[1], the auction reduced the share of premium seats sold at check-in from about 35% in year 1 to about 15% in year 2, suggesting that the airline prefers using the auction to allocate upgrades. About 7.5% of total premium capacity was allocated through the auction mechanism in the fifth quarter of data. 34% of total premium capacity was allocated through the check-in mechanism in the first quarter of data which dropped to 16% in the fifth quarter. Total allocated premium capacity increased slightly in the fifth quarter from 78% in the first quarter to 82% in the fifth quarter.

While it is tempting to claim that the auction eats into check-in upgrades more than full-price sales, such a claim cannot be made as there is a general increase in the number of premium seats sold across all flights due to changes in the revenue management system that took place around the time the check-in system was implemented. Furthermore, the auction system was implemented on top of the already existing check-in system, so it is hard to know how many passengers would substitute away from full-price premium tickets to the economy ticket with the auction option value. However, there is a very modest decrease in the seats sold at full price and with at check-in in Figure 6. This effect is not from the airline changing how often flights are flown as the effect can be seen in levels (Figure 6a) and in percentages (Figure 6b), where the effect is the most prominent. A similar effect can be observed in Figure 7 for premium revenues where the auction eats into check-in and full-price revenues over time.

[1] This is also observed in Figure 3

Figure 5: Share of Premium Seats Allocated Approaching Departure by Each Mechanism

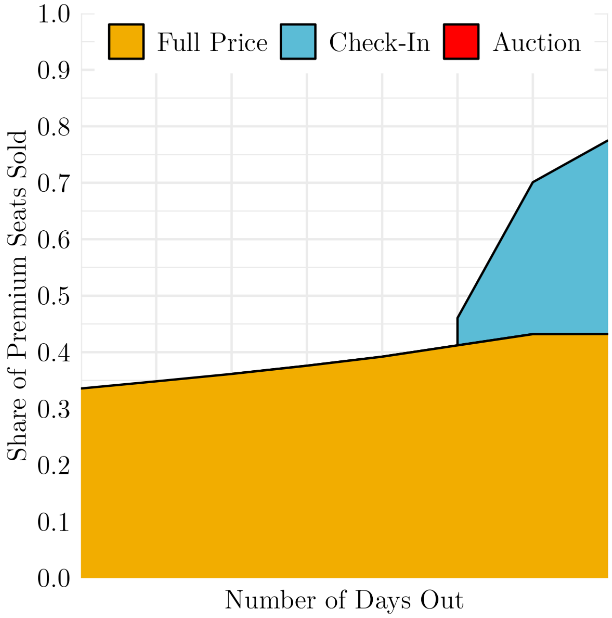

(a) Quarter 1

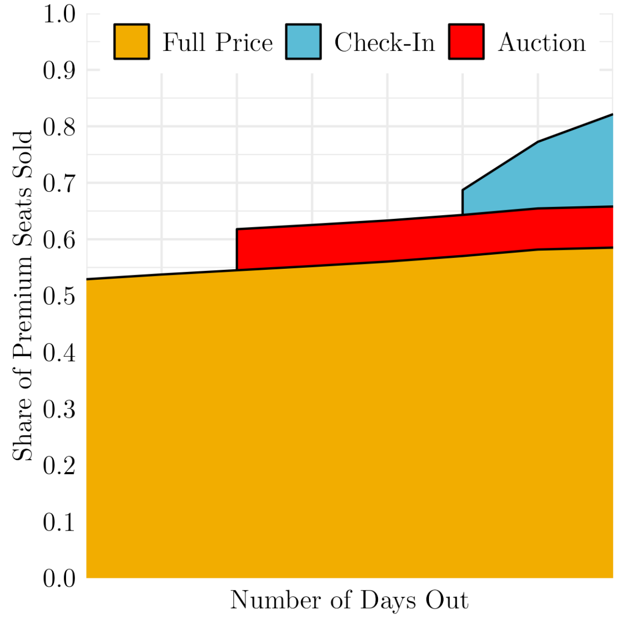

(b) Quarter 5

Notes:

Figure 6: Premium Seats Sold With Each Mechanism Over Time

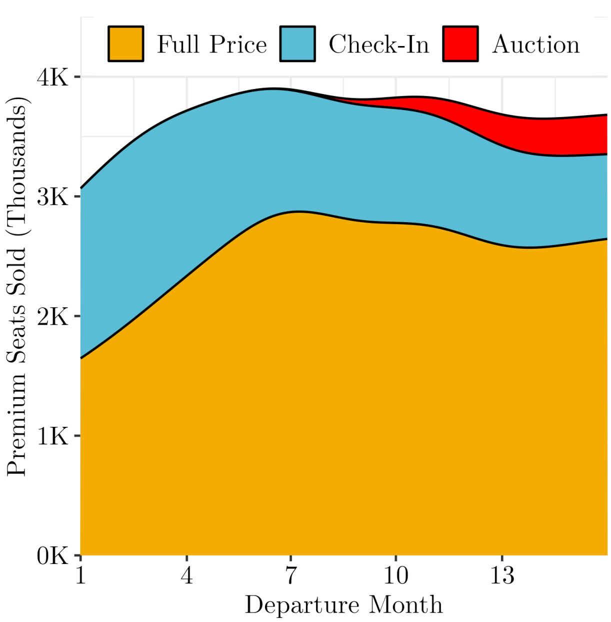

(a) Number of Premium Seats Sold by Purchase Channel

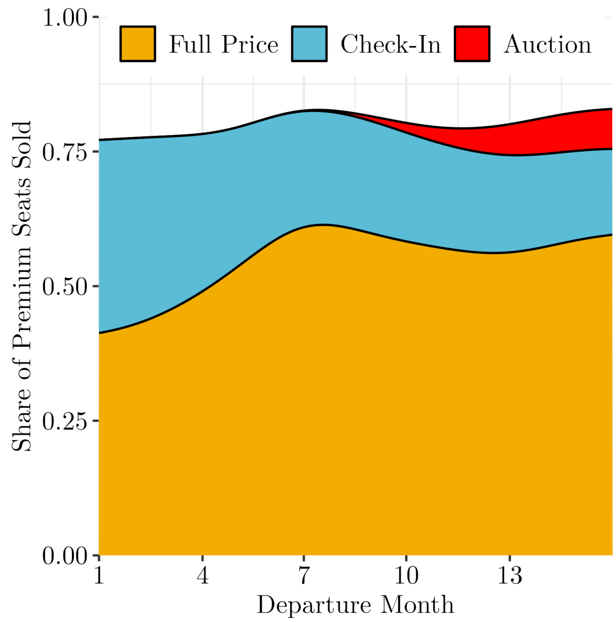

(b) Share of Premium Seats Sold by Purchase Channel

Notes:

4.3 Trade-offs and Revenue Implications

Figure 7: Impact of Each Mechanism on Premium Revenue Over Time

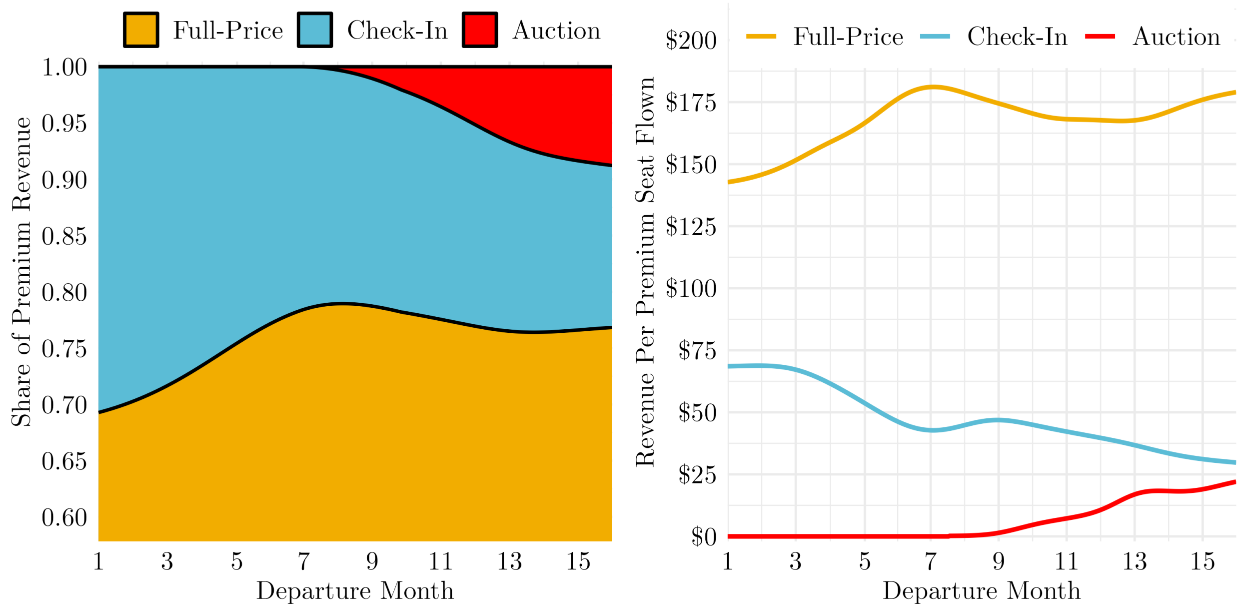

(a) Share of Premium Revenue by Mechanism

(b) Revenue Per Seat Flown

Notes:

The auction most notably eats into revenues from upgrades at check-in, which is consistent with Figure 6.

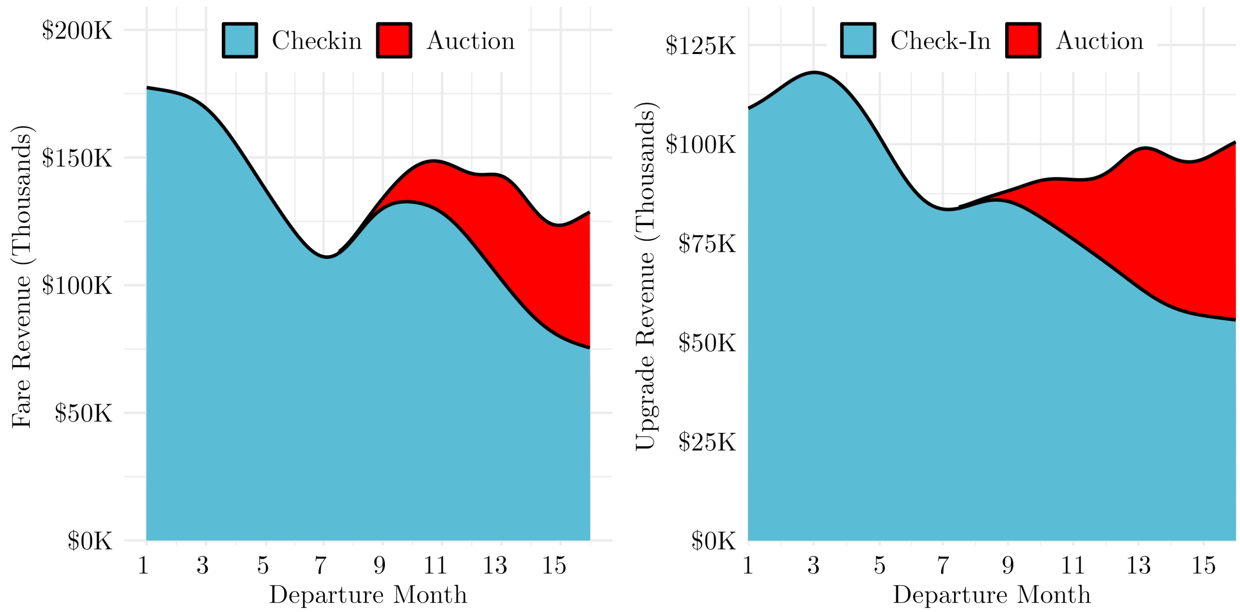

(a) Fare Revenue

(b) Upgrade Fee Revenue

Figure 8: The auction most notably eats into revenues from upgrades at check-in, which is consistent with Figure 6.

Table 7: Effect of Upgrades on Revenues Per Flight

	Revenue				
	Economy	Premium (Full)	Check-In	Auction	Total
	(1)	(2)	(3)	(4)	(5)
Variables					
Economy Sold	104.10***	0.74***	0.66***	0.07*	105.57***
	(1.90)	(0.17)	(0.07)	(0.03)	(1.97)
Premium Sold: All	208.88***	235.96***	1.55*	2.12***	448.51***
	(10.42)	(3.49)	(0.74)	(0.33)	(12.26)
Premium Sold: Auction	-106.37***	-234.43***	19.66***	218.76***	-102.37***
	(17.37)	(4.11)	(1.32)	(5.16)	(18.80)
Premium Sold: Check-In	-87.15***	-235.02***	165.40***	0.76*	-156.00***
	(12.10)	(3.92)	(3.43)	(0.34)	(12.81)
Fixed-effects					
Flight Number	Yes	Yes	Yes	Yes	Yes
Departure Date	Yes	Yes	Yes	Yes	Yes
Departure Hour	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	74,750	74,750	74,750	74,750	74,750
R^2	0.783	0.746	0.567	0.721	0.810
Within R^2	0.439	0.511	0.459	0.664	0.474

Clustered (Flight Number) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Notes: Estimated with the `feols()` function in the `fixest` R package [Bergé, 2018]. The unit of observation is at the flight level. The independent variables are the number of seats sold at departure acquired through each mechanism. The dependent variables for Columns 1 and 2 are the respective revenues attributable to fares from transactions that were never involved in an upgrade. Economy fares that are upgraded are removed from the total in Column 1, and, likewise, Column 2 only includes fares from full-price premium purchases. The dependent variables for Column 3 and 4 are the total revenues attributable to each upgrade mechanisms. This includes the base economy fare as well as the fee attributable to the purchase of the upgrade. Finally, the dependent variable in Column 5 is the total revenue of the flight i.e. it is the sum of the dependent variables in Columns 1-4.

Table 8: Effect of Upgrades on Revenues Per Flight

	Revenue				
	Economy (1)	Premium (Full) (2)	Check-In (3)	Auction (4)	Total (5)
Variables					
Economy Sold	107.41***	0.41	0.88***	0.15*	108.85***
	(2.14)	(0.21)	(0.09)	(0.07)	(2.23)
Premium Sold: All	231.80***	232.36***	1.17	4.09***	469.41***
	(14.42)	(3.70)	(1.12)	(0.80)	(16.25)
Premium Sold: Auction	-322.78***	-240.77***	21.69***	209.86***	-332.01***
	(25.11)	(5.13)	(1.63)	(5.43)	(25.72)
Premium Sold: Check-In	-90.16***	-235.27***	170.83***	7.34***	-147.26***
	(15.56)	(4.52)	(3.73)	(0.75)	(16.97)
Fixed-effects					
Flight Number	Yes	Yes	Yes	Yes	Yes
Departure Date	Yes	Yes	Yes	Yes	Yes
Departure Hour	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	30,413	30,413	30,413	30,413	30,413
R^2	0.774	0.747	0.575	0.679	0.800
Within R^2	0.423	0.489	0.474	0.578	0.447

Clustered (Flight Number) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Notes: Estimated with the `feols()` function in the `fixest` R package [Bergé, 2018]. The unit of observation is at the flight level. The independent variables are the number of seats sold at departure acquired through each mechanism. The dependent variables for Columns 1 and 2 are the respective revenues attributable to fares from transactions that were never involved in an upgrade. Economy fares that are upgraded are removed from the total in Column 1, and, likewise, Column 2 only includes fares from full-price premium purchases. The dependent variables for Column 3 and 4 are the total revenues attributable to each upgrade mechanisms. This includes the base economy fare as well as the fee attributable to the purchase of the upgrade. Finally, the dependent variable in Column 5 is the total revenue of the flight i.e. it is the sum of the dependent variables in Columns 1-4.

Table 9: Revenue Per Flight Regressions

	Revenue				
	Economy (1)	Premium (Full) (2)	Check-In (3)	Auction (4)	Total (5)
Variables					
Economy Sold	104.19***	0.85***	0.66***	0.10**	105.80***
	(1.90)	(0.17)	(0.07)	(0.04)	(1.97)
Premium Sold: Full	210.28***	237.38***	1.53*	1.52***	450.72***
	(10.47)	(3.53)	(0.75)	(0.31)	(12.32)
Premium Sold: Auction	100.17***	-1.21	21.18***	220.48***	340.62***
	(16.27)	(3.16)	(1.39)	(5.16)	(18.85)
Premium Sold: Check-In	123.86***	3.23	166.95***	2.46***	296.49***
	(12.55)	(2.37)	(3.56)	(0.40)	(14.25)
Fixed-effects					
Flight Number	Yes	Yes	Yes	Yes	Yes
Departure Date	Yes	Yes	Yes	Yes	Yes
Departure Hour	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	74,750	74,750	74,750	74,750	74,750
R^2	0.783	0.746	0.567	0.721	0.810
Within R^2	0.439	0.512	0.459	0.664	0.474

Clustered (Flight Number) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Notes: Estimated with the `feols()` function in the `fixest` R package [Bergé, 2018]. The unit of observation is at the flight level. The independent variables are the number of seats sold at departure acquired through each mechanism. The dependent variables for Columns 1 and 2 are the respective revenues attributable to fares from transactions that were never involved in an upgrade. Economy fares that are upgraded are removed from the total in Column 1, and, likewise, Column 2 only includes fares from full-price premium purchases. The dependent variables for Column 3 and 4 are the total revenues attributable to each upgrade mechanisms. This includes the base economy fare as well as the fee attributable to the purchase of the upgrade. Finally, the dependent variable in Column 5 is the total revenue of the flight i.e. it is the sum of the dependent variables in Columns 1-4.

4.4 Revenue Regression

$$\log(y_{nt} + 1) = \text{Auction}_{nt}\beta_{ATE} + \alpha_n + \lambda_t + \varepsilon_{nt} \tag{8}$$

where y_{nt} is a revenue variable for flight number n on departure date t, Auction_{nt} is an indicator for if flight number n had the auction on departure date t, β_{ATE} is the coefficient on the auction indicator which is the average treatment effect, α_n is a fixed-effect for flight number n, and λ_t is a fixed-effect for departure date t. Table **??** contains the average treatment effect on average economy revenue, premium revenue, and total revenue scaled per seat mile flown.

4.5 Behavioral Considerations

Consumers in our model are fully informed about the auction and have fully rational beliefs regarding auction and check-in upgrades. In reality, these assumptions might be too strong and consumer behavior may differ from that of the model. For example, the way the auction is presented to consumers might influence their bidding behavior or they might not be fully informed about the auction. We explore evidence for deviations from fully rational behavior in the rest of this section.

Information and Notifications

Our model assumes that all consumers are aware of the auction and check-in upgrade mechanisms. Although this might be a plausible assumption later in the sample, it is likely auction participation increased over time as consumers became aware of the auction. Because the auction is introduced in the middle of the sample, we can observe consumers adopting the auction over time and examine what influenced consumers to participate in the auction.

The most noticeable effect in our data on participation is through notifications to passengers about the auction as well as making bidding easier by increasing the channels in which passengers can place bids. When the auction was first introduced at the beginning of the seventh month of data, passengers could place bids through two channels: the airline's mobile app or the marketing page for the auction on the airline's website. Shortly after, the airline added another channel where bids could be placed on the confirmation page immediately after purchasing a ticket. The last two bid channels to be added during our sample period were a pre-travel email that is sent to passengers who have purchased an economy ticket on a qualifying flight and on the passenger's booking page which shows their booked flights.

Figure 9 shows how auction participation increased over time, when bids are submitted relative to departure, and how the introduction of new bid channels and notifications impacted participation. Even though participation was increasing steadily since the auction was introduced, the pre-travel email notification and bid channel approximately doubled the number of bids for all flights on the same departure date, as seen in Figure 9a. After the pre-travel email was introduced, the increasing trend in bids submitted through the existing channels (the mobile app, marketing page, and confirmation page) is dampened, suggesting that the pre-travel email did cannibalize some bids that would have been placed through existing channels. However, considering that the doubling of bids happens immediately after the introduction of the pre-travel email and the number of bids placed through this channel remains high, it is unlikely that the increase in bids is exclusively from cannibalization. Figure 9b shows that the vast majority of bids through the email channel are placed closer to departure when the frequency of pre-travel emails is highest.

Figure 9: Number of Bids for All Flights by Bid Channels

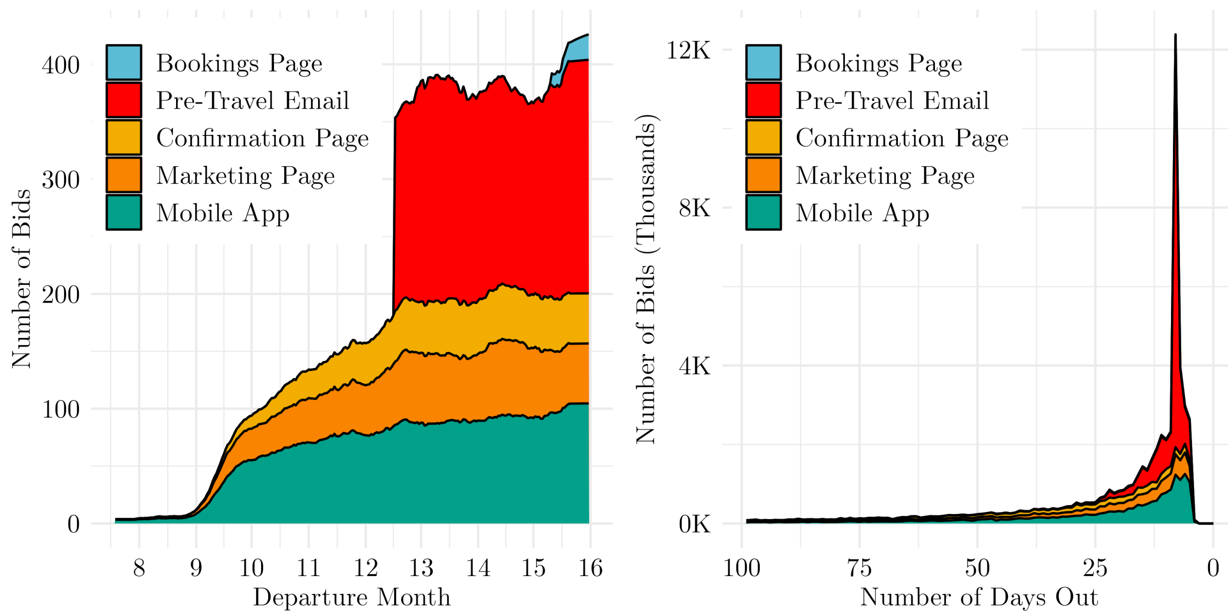

(a) Number of Bids by Departure Date (b) Number of Bids Approaching Departure

Notes: The number of bids in Figure 9a using a K-nearest neighbors Gaussian kernel regression using $K = 21$ and the rule-of-thumb bandwidth. The number of bids in Figure 9b are not smoothed.

An increase in auction participation can affect auction outcomes in a number of ways. First, in the canonical first-price sealed bid auction without endogenous entry, increased participation should increase bids by lessening the extent bidders shave down their valuations, consequentially increasing auction revenue. Second, the pre-travel email may have made consumers aware of the auction resulting in some flights receiving a positive number of bids when they otherwise would have received zero bids. The increased information should unambiguously increase auction revenue and bids.[2] Although this is related to the increased participation discussed in the previous point, it is different in that it is unclear if the bidding behavior of those made aware of the auction through the pre-travel email should be any different than those who were already aware of the auction. Lastly, related to the second point, if the passengers made aware of the auction from the pre-travel email are inherently "less sophisticated" consumers and submit lower, "irrational" bids, the average submitted bid could decrease with an ambiguous effect on average auction revenue.

We use the introduction of the pre-travel email to examine how increased auction participation from notifications affected bids and auction revenue. Figure 10 shows the effects of the emails on the average number size of submitted and accepted bids. Across all flights, the implementation of the pre-travel email resulted 0.75 more submitted bids and 0.5 accepted bids as seen in Figures 10a and 10b, respectively. These effects are larger for flight numbers that were receiving more bids on average before to the email, suggesting that the email was notifying uninformed passengers of the auction or reminding forgetful passengers to bid. The increase in participation and consequentially accepted bids also increases bid revenue by about \$50 across all flights highlighted in Figure 10c. While the effect on the number of submitted and accepted bids appears to differ across the distribution of flight numbers, this does not appear to be the case with average bid revenue. The reason for this is because the average submitted bid decreases slightly when the pre-travel email is introduced, as seen in Figure 10d. In fact, the average bid decreases the most for the top 50% of flight numbers with the most submitted bids before the email. As these are likely flights where the upgrade is valued more by consumers, this suggesting that the email attracted bidders who placed lower bids further suggesting that these bidders either value the upgrade less or are less sophisticated bidders.

Table 10 shows the effects discussed above estimated with a linear regression in columns 1 through 4. The estimates similar both qualitatively and quantitatively. Column 5 shows the difference in average submitted bids through each bid channel with the omitted group being bids placed on the confirmation page immediately after purchasing tickets. Compared to bids placed on the confirmation page, bids placed through the marketing page[3] are \$3.50 higher on average. However, bids placed through the mobile app and the pre-travel are on average \$6.15 and \$4.73 lower, respectively. Bids placed through the confirmation page immediately after purchase are likely submitted by consumers behaving most similar to the consumers in our model.

Framing from Slider Presentation

[2]We define the average bid when no bids are submitted to be zero.

[3]The marketing page is the page on the airline's website that advertises the upgrade auction.

Figure 10: Impact of Pre-Travel Email on Participation, Bidding, and Auction Revenue

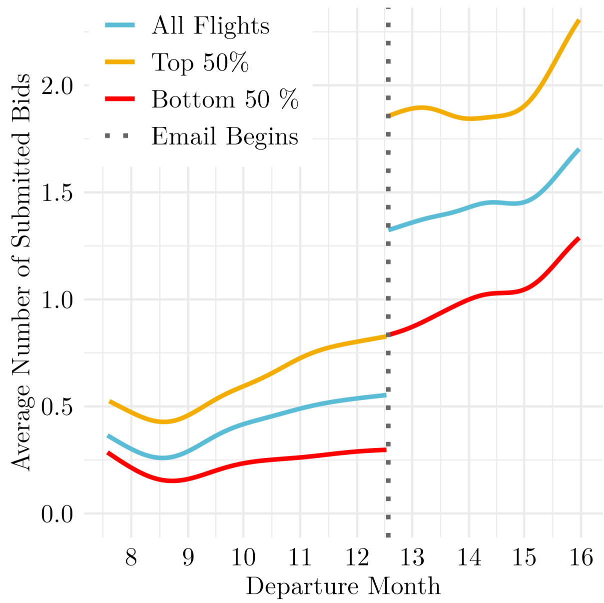

(a) Average Number of Bids Per Flight

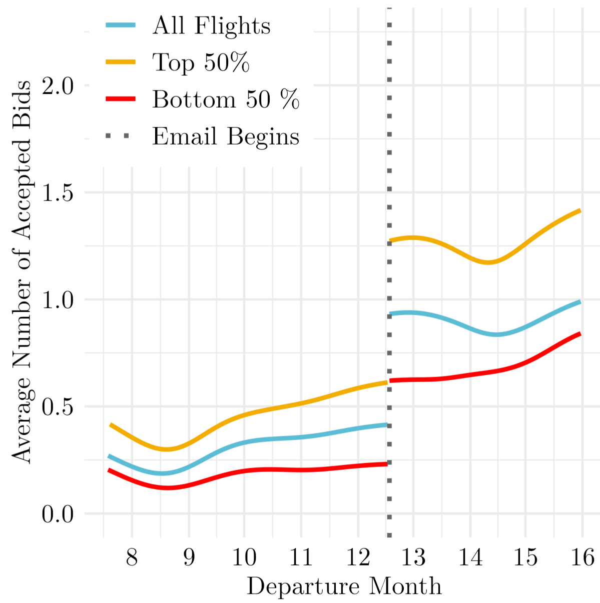

(b) Average Accepted Bids Per Flight

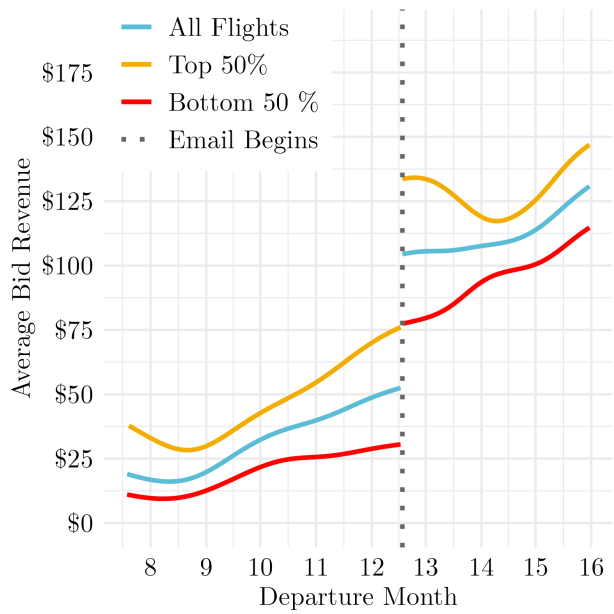

(c) Average Auction Revenue Per Flight

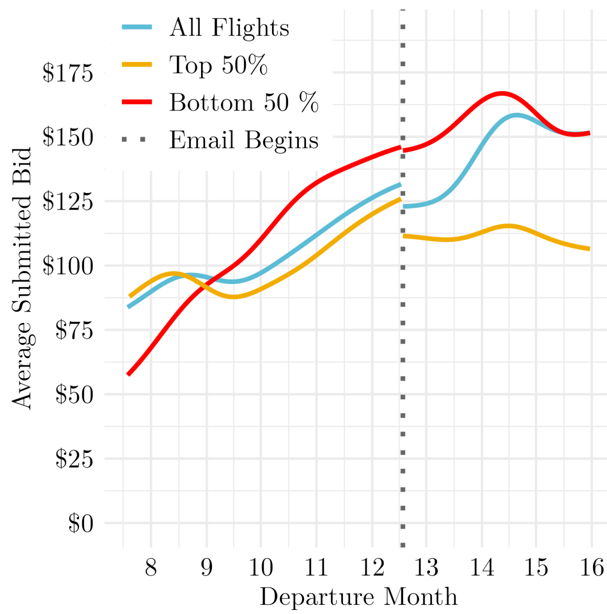

(d) Average Submitted Bid

Notes: Each variable is measured at the flight level. The top and bottom 50% refers to the share of flight numbers sorted by their average number of submitted bids in the period before the pre-travel email. Averages are computed by smoothing flight level variables with a Gaussian kernel regression and the rule-of-thumb bandwidth. Smoothing is performed separately before and after the pre-travel email introduction.

Table 10: Pre-Travel Email Affect On Flight Auctions

	Number of Submitted Bids	Number of Accepted Bids	Bid Revenue	Submitted Bid	
	(1)	(2)	(3)	(4)	(5)
Variables					
Post Email	0.7519***	0.5078***	48.33***	-0.0341	
	(0.0336)	(0.0258)	(3.653)	(0.9006)	
Marketing Page					3.503***
					(0.8216)
Mobile App					-6.150***
					(0.7364)
Pre-Travel Email					-4.737***
					(0.6605)
Fixed-effects					
Flight Number	Yes	Yes	Yes	Yes	Yes
Flight Month	Yes	Yes	Yes	Yes	Yes
Flight Day of Week	Yes	Yes	Yes	Yes	Yes
Departure Hour	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	74,757	74,757	74,757	74,933	74,366
R^2	0.21897	0.18224	0.15856	0.89572	0.89717

Clustered (Flight Number) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Notes: Estimated with the `feols()` function in the `fixest` R package [Bergé, 2018]. The excluded group for column 5 is bids placed at the confirmation page immediately after purchasing a ticket. Bids placed through the bookings page are excluded from the sample in column 5 as the bid channel was implemented near the end of the sample and only resulted in 567 bids.

Because passengers place bids with a slider as shown in Figure 2, it is possible that bidding is affected by the way the slider is presented. The slider affects bidding a few different ways. First, it induces endogenous entry as passengers who are not willing to bid the minimum do not enter the auction (or effectively have a bid of zero). Secondly, the slider causes bids to be discrete as the slider increases in increments of $5. Both of these are entirely consistent with the consumers in our model. However, the initial position of the slider can influence bidding behavior whether through rational information extraction or behavioral framing effects. We explore to what extend there is evidence of framing in the data.

Figure 11 shows the distribution of the "normalized" bids and slider starting position. The normalization of these variables transforms them into a percentage between the slider

Figure 11: Evidence of Framing Effect in Bid Distributions

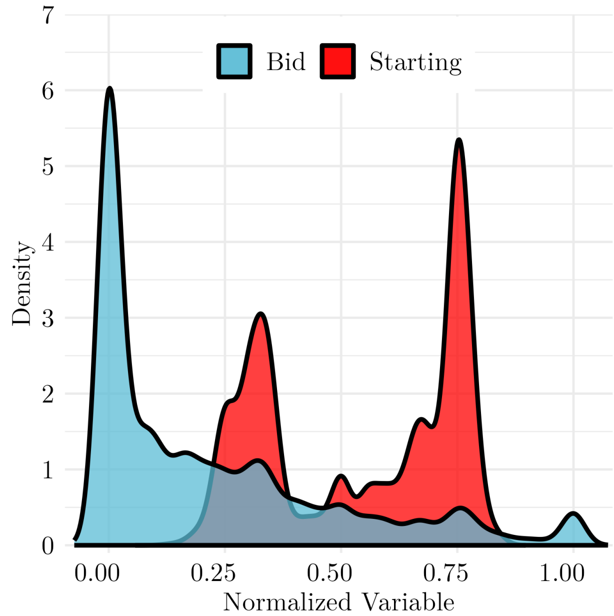

(a) Distribution of normalized bid and starting position

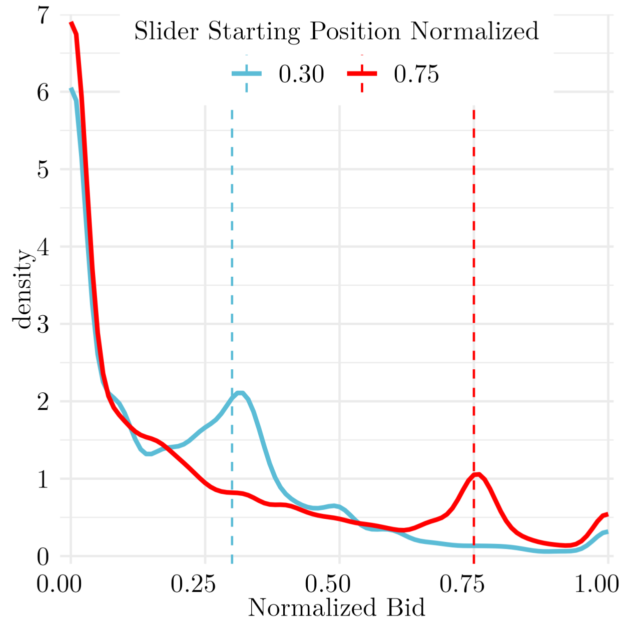

(b) Distribution of normalized bid conditional on normalized starting position

Notes: The densities in a are estimated with the `density()` function from the `stats` package in base R [R Core Team, 2021]. The conditional densities in b are estimated with the `cde()` function in the `hdrcde` R package [Hyndman et al., 2021] using the default bandwidths.

minimum and maximum values. The most common bid was the minimum and the most common initialize positions were at the 0.75 and 0.30 positions on the slider. There are noticeable humps in the distribution of normalized bids around the 0.30 and 0.75 position. To check if these humps in the normalized bid distribution correspond to the the initial slider positions, Figure 11b shows the conditional distribution of the normalized bid given different normalized starting positions of 0.30 and 0.75. For each initial position, the second most common bid after the slider minimum is around the initial slider position, suggesting that consumers condition their bid on the starting position.

The humps in the conditional distributions suggest that this is a framing effect rather than passengers (rationally) extracting information from the initial position of the slider. Assuming that a higher initial position informs passengers that the auction is more competitive and a higher bid is necessary to obtain an upgrade, one would expect the initial slider to move the density of the conditional bid distribution to the right, increasing the average of the conditional bid distribution. That is, the CDF conditional on a starting position of 0.75 should first order stochastically dominate the CDF conditional on a starting position of 0.30. However, that is not what we observe (see Figure 13 in the Appendix) and instead see probability density pooling around the starting position with little change to the rest of the distribution.

This framing effect can be seen in Table 11 where a a 1% increase in the starting position of the slider is associated with a 0.11% to 0.15% depending upon the specification. In column 1, a one dollar increase in the slider starting value is associated with an 11.5 cent increase in submitted bids conditional on the slider minimum and maximum values. However, notice that a one dollar increase in the slider minimum value is associated with a 87.2 cent increase in submitted bids. The fact that this effect is less than one should be expected because increasing the slider minimum should result in a smaller increase in submitted bids if the airline is setting the slider minimum optimally. In column 2, a one percentage point increase in the normalized starting position is associated with a 14.8 cent increase in submitted bids conditional on the slider minimum and maximum values. Lastly, column 3 shows that a one percentage point increase in the normalized starting position is associated with a 0.152 percentage point increase in the normalized bid.

Table 11: Bidding Regressions: Framing

	Bid		Normalized Bid
	(1)	(2)	(3)
Variables			
Slider Minimum	0.8857***	0.9427***	
	(0.0332)	(0.0279)	
Slider Maximum	0.0113	0.0575*	
	(0.0254)	(0.0280)	
Slider Starting	0.1176***		
	(0.0209)		
Normalized Starting Position		0.1410***	0.1769***
		(0.0194)	(0.0134)
Fixed-effects			
Flight Number	Yes	Yes	Yes
Departure Date	Yes	Yes	Yes
Departure Day of Week	Yes	Yes	
Fit statistics			
Observations	72,467	72,467	72,467
R^2	0.90782	0.90772	0.06592

Clustered (Flight Number) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Notes: Estimated with the `feols()` function in the `fixest` R package [Bergé, 2018].

Table 12: Bidding Regressions: Framing

	Bid		Normalized Bid
	(1)	(2)	(3)
Variables			
Slider Minimum	0.8728***	0.9327***	
	(0.0347)	(0.0285)	
Slider Maximum	-0.0013	0.0476	
	(0.0267)	(0.0288)	
Slider Starting	0.1241***		
	(0.0226)		
Economy Fare Paid	0.0348***	0.0346***	
	(0.0037)	(0.0037)	
Normalized Starting Position		0.1489***	0.1810***
		(0.0207)	(0.0140)
log(Economy Fare Paid)			0.0441***
			(0.0036)
Fixed-effects			
Flight Number	Yes	Yes	Yes
Departure Date	Yes	Yes	Yes
Departure Day of Week	Yes	Yes	
Fit statistics			
Observations	67,521	67,521	67,521
R^2	0.90898	0.90886	0.07256

Clustered (Flight Number) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Notes: Estimated with the `feols()` function in the `fixest` R package [Bergé, 2018].

5 Conclusions

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 and Maria. F. Gabrielli. An empirical analysis of competitive nonlinear pricing. *International Journal of Industrial Organization*, 68:102538, 2019.

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.

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.

Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–57, March 2015. doi: 10.1257/aer.20130848. URL https://www.aeaweb.org/articles?id=10.1257/aer.20130848.

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.

G.S. Crawford and Matthew Shum. Monopoly quality degradation in cable television. *Journal of Law and Economics*, 49(1):181–209, 2006.

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, 2018. 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.

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.

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.

Igal Hendel and Aviv Nevo. Measuring the implications of sales and consumer inventory behavior. *Econometrica*, 74(6):1637–1673, 2006.

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*, 139(2):1149–1199, 09 2023. ISSN 0033-5533. doi: 10.1093/qje/qjad051. URL https://doi.org/10.1093/qje/qjad051.

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

Phillip Leslie. Price discrimination in broadway theater. *RAND Journal of Economics*, 35 (3):520–541, 2004.

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.

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

Brian McManus. Nonlinear pricing in an oligopoly market: The case of specialty coffee. *Rand Journal of Economics*, 38(2):512–532, 2007.

Julie Holland Mortimer. Price discrimination, copyright law, and technological innovation: Evidence from the introduction of dvds. *The Quarterly Journal of Economics*, 122(3): 1307–1350, 2007.

Harikesh Nair. Intertemporal price discrimination with forward-looking consumers: Application to the us market for console video-games. *Quantitative Marketing and Economics*, 5(3):239–292, 2007.

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.

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.

Hal R. Varian. Price discrimination and social welfare. *The American Economic Review*, 75 (4):870–875, 1985. ISSN 00028282. URL http://www.jstor.org/stable/1821366.

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.

Computational References

Laurent Bergé. Efficient estimation of maximum likelihood models with multiple fixed-effects: the R package FENmlm. *CREA Discussion Papers*, (13), 2018. URL https://cran.r-project.org/web/packages/fixest/index.html.

Rob J Hyndman, Jochen Einbeck, and Matthew P Wand. *hdrcde: Highest Density Regions and Conditional Density Estimation*, 2021. URL https://pkg.robjhyndman.com/hdrcde/. R package version 3.4.

R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2021. URL https://www.R-project.org/.

A Additional Figures and Tables

Figure 12: Distribution of Capacity States Before Auction

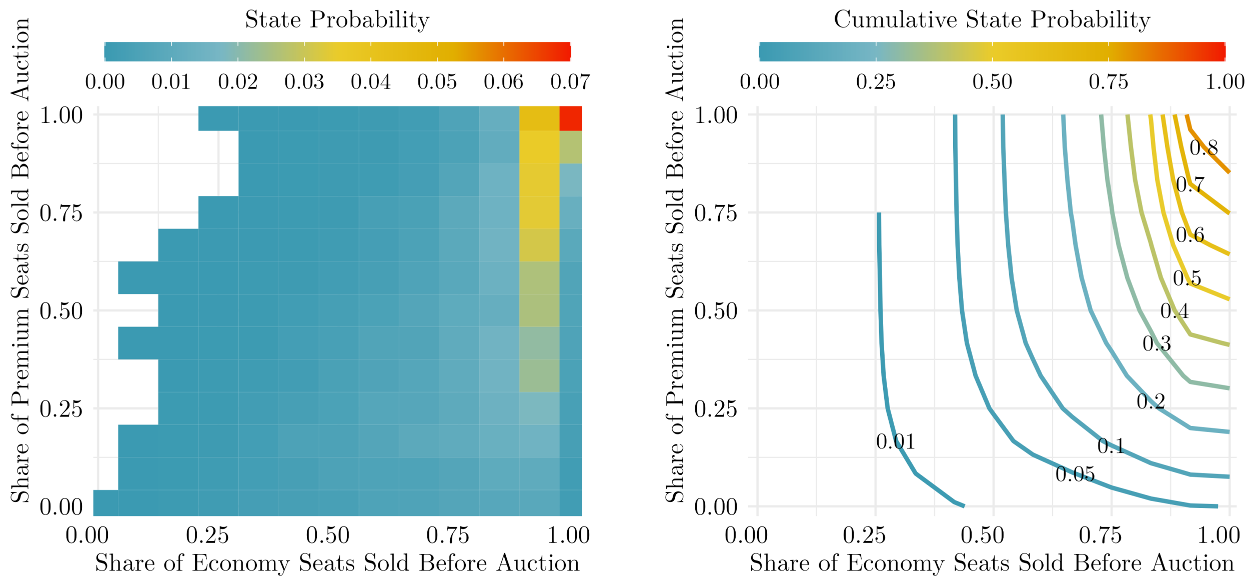

(a) Joint Distribution Heat-Map (b) Joint Cumulative Distribution Contours

Notes: A "state" is an ordered pair of the share of remaining seats (capacity) in each cabin the day before the auction. A tile on the heat-map represents one possible state. The economy and premium remaining capacity shares are binned by rounding to the nearest $1/12^{\text{th}}$.

Table 13: Effect of Upgrades on Revenue Per Seat Flown

	Revenue Per Seat Flown				
	Economy (1)	Premium (Full) (2)	Check-In (3)	Auction (4)	Total (5)
Variables					
Economy Sold	0.38***	0.01	0.05***	0.00	0.33***
	(0.01)	(0.02)	(0.01)	(0.00)	(0.01)
Premium Sold: All	2.75***	18.74***	0.02	0.13***	4.28***
	(0.11)	(0.38)	(0.07)	(0.03)	(0.13)
Premium Sold: Auction	-1.45***	-19.24***	1.62***	18.02***	-1.37***
	(0.13)	(0.36)	(0.11)	(0.42)	(0.13)
Premium Sold: Check-In	-1.11***	-19.60***	13.64***	0.04	-1.59***
	(0.10)	(0.34)	(0.28)	(0.03)	(0.10)
Fixed-effects					
Flight Number	Yes	Yes	Yes	Yes	Yes
Departure Date	Yes	Yes	Yes	Yes	Yes
Departure Hour	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	74,750	74,750	74,750	74,750	74,750
R^2	0.681	0.712	0.564	0.725	0.719
Within R^2	0.233	0.490	0.456	0.668	0.289

Clustered (Flight Number) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Notes: Estimated with the `feols()` function in the `fixest` R package [Bergé, 2018]. The unit of observation is at the flight level. The independent variables are the number of seats sold at departure acquired through each mechanism. The dependent variables for Columns 1 and 2 are the respective revenues per seat flown in each cabin attributable to transacted fares that were never involved in an upgrade. Economy fares that are upgraded are removed from the total in Column 1, and, likewise, Column 2 only includes fares from full-price premium purchases. The dependent variables for Column 3 and 4 are the total revenues attributable to each upgrade mechanisms per seat flown in the premium cabin. This includes the base economy fare as well as the fee attributable to the purchase of the upgrade. Finally, the dependent variable in Column 5 is the total revenue per seat flown of the flight i.e. it is the sum of the dependent variables in Columns 1-4.

Table 14: Revenue Per Seat Flown Regressions

	Revenue Per Seat Flown				
	Economy (1)	Premium (Full) (2)	Check-In (3)	Auction (4)	Total (5)
Variables					
Economy Sold	0.39***	0.02	0.05***	0.01*	0.33***
	(0.01)	(0.02)	(0.01)	(0.00)	(0.01)
Premium Sold: Full-Price	2.74***	18.85***	0.01	0.08**	4.28***
	(0.11)	(0.39)	(0.07)	(0.03)	(0.13)
Premium Sold: Auction	1.26***	-0.71*	1.64***	18.11***	2.85***
	(0.14)	(0.32)	(0.13)	(0.42)	(0.16)
Premium Sold: Check-In	1.64***	-0.68*	13.65***	0.13***	2.71***
	(0.14)	(0.33)	(0.30)	(0.04)	(0.16)
Fixed-effects					
Flight Number	Yes	Yes	Yes	Yes	Yes
Departure Date	Yes	Yes	Yes	Yes	Yes
Departure Hour	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	74,750	74,750	74,750	74,750	74,750
R^2	0.681	0.713	0.564	0.725	0.719
Within R^2	0.233	0.491	0.456	0.668	0.288

Clustered (Flight Number) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Notes: Estimated with the `feols()` function in the `fixest` R package [Bergé, 2018]. The unit of observation is at the flight level. The independent variables are the number of seats sold at departure acquired through each mechanism. The dependent variables for Columns 1 and 2 are the respective revenues per seat flown in each cabin attributable to transacted fares that were never involved in an upgrade. Economy fares that are upgraded are removed from the total in Column 1, and, likewise, Column 2 only includes fares from full-price premium purchases. The dependent variables for Column 3 and 4 are the total revenues attributable to each upgrade mechanisms per seat flown in the premium cabin. This includes the base economy fare as well as the fee attributable to the purchase of the upgrade. Finally, the dependent variable in Column 5 is the total revenue per seat flown of the flight i.e. it is the sum of the dependent variables in Columns 1-4.

Figure 13: Evidence of Framing Effect in Bid CDF

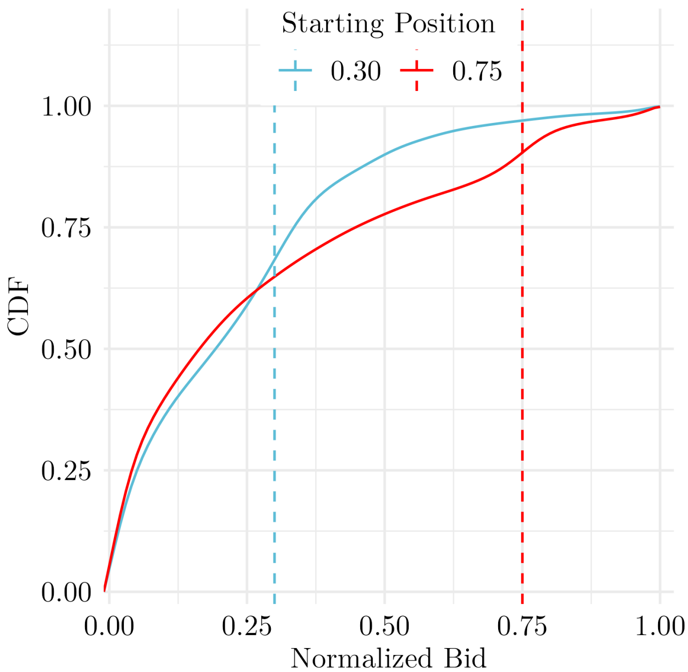

Notes: