Welfare Consequences of Upgrades: Evidence from the Airline Industry

Domenico Fabrizi*

UCLA

November 27, 2023

[FOR THE MOST RECENT VERSION CLICK HERE]

Abstract

Using upgrades—fees that customers pay to access a premium quality product after the purchase of a regular one—can significantly affect consumer welfare. On the one hand, regular consumers benefit from the ability to access a monopolist's higher-quality goods at a discounted price. On the other hand, the monopolist will seek to capture surplus from these gains from trade by offering a different price menu for all goods. Whether consumer welfare rises or falls when a firm introduces upgrades depends on the relative magnitudes of these two effects. The aim of this research is to disentangle the two effects in the context of an international airline that offers economy class passengers the option to pay an additional fee to upgrade to business class. I develop and estimate a model of airline pricing to assess the effects of such upgrades via counterfactual simulations. I show that the upgrade option improves the allocation of passengers across cabins, which leads to an increase in consumer and producer welfare of 1.5% and 2%, respectively.

*Department of Economics, University of California, Los Angeles (email: domenicofabrizi@g.ucla.edu).
I thank an anonymous airline for the data used in this study. I am indebted to John Asker, Simon Board, Martin Hackmann, and, in particular, Will Rafey for their guidance and advice. I also thank Akira Ishide, Victoria Barone, Joaquin Serrano, Francesco Gabriele, Alvaro Boitier, Lorenzo Cattivelli, Chandni Raja, Andrew Heinzman, Manu Navjeevan, Daniel Perez, Calvin Kuo, Jesper Böjeryd, and numerous UCLA proseminar participants. All errors are my own.

1 Introduction

Offering upgrades is a popular practice in industries characterized by multiproduct offerings and limited capacity. Airlines, freight transport providers, car rental companies and hotels offer products of various quality and allow customers to access premium-quality products through either retail sales or upgrades. Upgrades can generate up to 10% and 8% of revenues in the car rental (AllianzTravel) and the airline (OAG) industries; however, they have received little attention in the empirical economics literature, because of lack of accessible data.

Firms complement the retail sales channel with upgrades to sequentially price discriminate and manage limited inventory. However, the effects on total welfare of the introduction of this new sales channel are ambiguous. On the one hand, the availability of an upgrade option provides the seller with extra flexibility in selling products, increasing producer surplus. In particular, through upgrades, firms can find customers who value premium products below the listed price among those who already hold a regular product. According to theoretical research, this may alleviate allocative inefficiencies by reducing mismatch between supply and demand (Gallego and Stefanescu, 2009; Cui, 2017). Moreover, existing evidence suggests that airlines' revenues could increase after the introduction of upgrade options (Cui et al., 2019). On the other hand, the effect on customers is less straightforward. Upgrades increase the welfare of consumers who can enjoy premium products at discounted prices. However, introducing an upgrade option modifies the seller's pricing problem, potentially influencing prices for retail sales, with ambiguous effects on consumer welfare. Existing theoretical studies (Varian, 1985; Aguirre et al., 2010; Bergmann et al., 2015) do not consider sequential price discrimination strategies featuring a dynamic multiproduct offering and limited inventory, thus leaving open the question of what the effects of upgrade options are on consumers.

This paper studies how introducing upgrades affects welfare in the context of the airline industry. In particular, I analyze new proprietary data from an international airline that employs upgrades to allocate premium-cabin seats.[1] I find both that upgrades are a relevant sales channel for premium products and that the airline employs them for

[1] The identity of the airline is confidential. In terms of size, in 2018, the airline studied in this paper had revenues between \$10 and \$20B. As a benchmark, American Airlines made \$43B (macrotrends).

price discrimination. I then estimate a structural model to quantify the effect of the upgrade option on welfare. My results indicate that, on average, both consumers and the monopolist benefit from the upgrade option.

There are at least three compelling reasons to study ticket upgrades in the airline industry. First, the airline sector is one of the world's largest industries, contributing approximately 5% to global GDP (IATA, 2018). Second, ticket upgrades are extensively used in the industry, with upgrade programs being implemented by all major international airlines (see, for example, the websites of United Airlines, Delta Airlines, American Airlines, China Southern Airlines, Lufthansa and KLM). Third, the airline industry has low profit margins, as margins average around 1.2% and profits per passenger equal $2.25 (IATA, 2023): this makes revenues generated from upgrades a vital contributor to the overall profitability of the industry (McKinsey, 2019).

In this paper, I use a novel transaction-level dataset obtained from the revenue management department of an international airline. The dataset contains all 268,035 ticket purchases made over 1,600 flights, connecting two monopolistic routes, during the period between September and December 2018. For each ticket in the dataset, I have access to information such as the final price paid, timing of purchase and cabin class (economy, premium economy or business class). Notably, the dataset shows whether customers in lower class cabins upgraded to premium cabins and, if so, at which price.

This dataset enables me to find new evidence about upgrades and validate well-established facts about retail sales in the airline industry. Regarding upgrades, I am able to establish two significant findings. First, upgrades serve as a valuable sales channel, with the airline utilizing them to sell 16% of their premium seat tickets, thereby generating 2% of total revenues. Second, the airline uses upgrades to sell higher quality seats at a discounted price. On average, customers who upgrade to premium cabins save 15% off the retail price of seats in those cabins. This is in line with the theory of sequential price discrimination, whereby the firm first selects demand through retail sales and then offers the high-quality product again at lower prices. The retail pricing data also exhibit characteristics consistent with those described in the economics literature on airline pricing. In line with Williams (2022) and Lazarev (2013), the airline engages in dynamic pricing to implement intertemporal price discrimination among customers. Those who purchase their tickets early appear to have a lower willingness to pay (WTP) for travel than those who purchase at the last minute. Furthermore, in line with Aryal et al. (2023), the airline engages in intratemporal price discrimination by offering products of different qualities at various prices to segment customers based on their value for comfort.

To measure the changes in consumer and producer welfare resulting from the upgrade sales channel, I construct a structural model of supply and demand, building on Aryal et al. (2023). On the supply side, I approximate the airline decision process by a two-period problem with capacity constraints, where the airline chooses prices for two products: economy and business class seats. In both periods, the airline sells economy and business class seats through the retail channel. In the last period, given availability in the premium cabin, the airline sells upgrades to customers holding an economy class ticket from the initial period. As a result, the airline's dynamic programming problem encompasses multiple products and upgrades. The model allows me to capture crucial aspects of airline price discrimination strategies, including intertemporal (dynamic pricing), intratemporal (regular vs premium cabin) and sequential (upgrades) price discrimination.

On the demand side, customers randomly enter the airline ticket market and, upon arrival, make decisions to purchase an economy or a business class ticket, or not to fly. If customers buy an economy class ticket in $t = 0$, they have the option to upgrade to business class in the last period, $t = 1$. When modeling customer's upgrade option, I make two assumptions. Firstly, I impose that customers do not anticipate the possibility of an upgrade. While this assumption simplifies the problem solution and aligns with the practice in much of the prior literature,[2] it rules out strategic thinking on the part of consumers. Additionally, I allow customers to be inattentive. Specifically, the airline sends a notification (e.g., an email) to all economy class ticket holders two days before departure, but only a fraction of them actually pay attention to the notification itself. Allowing for inattention enhances the model's credibility in two ways. First, it mirrors how firms in the airline industry offer upgrades (as reported by Tripadvisor; see also Appendix A.1 for an example). Second, it aligns with customers' patterns of response to notifications in the travel industry, where only 20% of customers open travel-related emails (as reported by Constant Contact and campaignmonitor.com).

[2] For example, Williams (2022) and Aryal et al. (2023) implement a similar assumption; one relevant exception is Lazerev (2013), who allows consumers to strategically time their initial ticket purchase.

To account for price differences across flights and over time, I specify a model with random preference coefficients. Each simulated flight at each time is characterized by an idiosyncratic demand shock, which characterizes flight-time specific distributions of WTP and value for comfort. The airline sets optimal prices based on these demand shocks. Variation in demand shocks across flights and over time, then, leads to variation in simulated prices and sales. By assuming optimal pricing alongside demand-side restrictions, I am able to identify customers' preferences and arrival rate over time. I estimate the model at the route level using the simulated method of moments, where I match the distribution of prices and ticket sales.[3] The estimated parameters allow the model to approximate the observed distributions of prices and quantities.

Using these estimates, I conduct counterfactual simulations to examine how the upgrade option affects pricing decisions and welfare. First, I discuss the impact of upgrades on pricing decisions and their interaction with capacity constraints conditional on the demand shock. On the one hand, irrespective of capacity constraints, the airline uses upgrades to sequentially price discriminate among initial-period customers, thereby affecting initial-period prices. In particular, since the airline anticipates that some initial-period economy class ticket holders will upgrade to business class, the opportunity cost of raising initial-period premium-cabin prices decreases. However, how the upgrade option interacts with capacity constraints and how it affects the airline pricing problem depend on the demand shock. In the case of a low demand shock in the initial period, the airline experiences few sales in both cabins, and it uses the upgrade option only to sequentially price discriminate among initial-period customers: this alone does not affect last-period prices. On low-value flights, then, the upgrade option increases welfare by enabling the airline to fill empty seats in the premium cabin and granting customers access to this cabin at a discount. In contrast, in the case of large, positive initial-period demand shocks, the probability of business class selling out in the initial period is large and the airline potentially misses out on high-value sales from high-value travelers in the last period. In such a scenario, the airline uses the upgrade option as a costless way to reduce this probability. By anticipating that high-value customers who have bought economy class tickets in the initial period can upgrade in the final period, the airline further increases initial-period business class prices as a way to restrict the number of its retail sales in the

[3] Focusing on route-level demand is consistent with practices in the industry (see IATA).

initial period. Furthermore, beyond distorting initial-period prices, the upgrade option also affects prices in the last period due to its interaction with capacity constraints. In fact, upgrading customers compete with retail buyers for business class in the final period. On high-value flights, then, the upgrade option increases welfare by reducing the number of early sellouts in the premium cabin and improving the efficiency of seat allocation over time.

Second, I quantify the welfare effects of the upgrade option for both travelers and the airline. From a welfare standpoint, the availability of the upgrade option has a positive impact on both producer and consumer surplus, thereby increasing efficiency. Overall, producer surplus increases by an average of $1,553 (2%) per flight, which is significant in an industry with a 1.2% profit margin. Overall, consumer surplus increases by an average of $1,227 (1.5%) per flight, which is equivalent to consumer gains associated to a $6.50 subsidy to all boarding travelers, when the airline does not implement the upgrade option. The increase in producer surplus is driven by the revenue generated from the upgrade fees, which compensates for the fewer sales resulting from higher prices. Concurrently, the main driver of the increase in consumer surplus is the utility gain experienced by customers who choose to upgrade, thereby accessing business class at a discount. This compensates the loss in utility due to the higher retail prices of the premium product.

Literature review

This paper builds on three domains of literature. First, a recent literature in operations management that focuses on upgrading. Most papers study how sellers can benefit from upgrades from both the theoretical (Gallego and Stefanescu, 2009; Cui et al., 2018; Cui and Shin, 2018) and empirical perspectives (Cui et al., 2019; Yilmaz et al., 2017). In particular, Cui et al. (2019) study the effects of specific upgrade products known as add-ons, which require the purchase of a regular product before the premium product can be accessed, in a quasi-experiment within the airline industry. I extend this literature in at least two ways. One, none of these studies empirically quantify the consumer welfare changes resulting from upgrades. Two, I incorporate the possibility of customer inattention to the upgrade option, modeling it as a behavioral parameter following an approach similar to Gabaix (2019).

Second, this paper also extends the empirical literature studying how price discrimination affects welfare. Most empirical papers focus on either intratemporal (Leslie, 2004; Crawford and Shum, 2007) or intertemporal price discrimination (Lazarev, 2013; Cho et al., 2018). Only Chandra (2020) and Aryal et al. (2023) consider both aspects. I contribute to this line of research by developing an empirical model that allows me to consider upgrades as an additional sales channel through which the airline engages in sequential price discrimination.

Finally, this paper builds on recent airline industry studies, including Dana et al. (2022), Aryal et al. (2023), Williams (2022), Lazarev (2013), and Li et al. (2014). Williams (2022) and Lazarev (2013) analyze the effects of dynamic pricing on consumers; Li et al. (2014) investigate customers' incentives to delay ticket purchases, and Aryal et al. (2023) examine a monopolist selling vertically differentiated products with dynamic pricing. Meanwhile, Dana et al. (2022) consider a dynamic pricing inventory control problem. However, these papers do not take into account the role of the upgrade option in sequential price discrimination and inventory management, overlooking important factors that influence pricing decisions.

In particular, I build on the methods used by Williams (2022) and Aryal et al. (2023) in modeling and estimation. Williams (2022) considers a single product monopolistic airline engaging in dynamic pricing. Aryal et al. (2023) expand this framework, by investigating a firm combining dynamic pricing with vertical product differentiation. My paper further enriches the setting described by Aryal et al. (2023) by allowing the airline to sell upgrades. In terms of demand and supply modeling, my framework introduces a novel decision problem on both the demand side—as customers decide whether to upgrade—and the supply side, as the firm optimally chooses the upgrade price. In terms of identification strategy, I adopt the approaches of both Williams (2022) and Aryal et al. (2023) by leveraging supply-side constraints to argue separate identification of arrival rates and customer preferences. As for estimation, I follow Aryal et al. (2023) by implementing the simulated method of moments to estimate a random coefficient demand model. In terms of findings, my paper is closely related to that of Williams (2022), who demonstrates that his airline uses dynamic pricing to intertemporally price discriminate and to manage inventory, by reserving capacity for high-value customers in later periods. Similarly, my results indicate that upgrades not only work as a way to price discriminate but also, conditional on the size of the demand shock, can serve as a tool to improve inventory allocation over time.

2 Data and descriptive evidence

This project uses a new transaction-level dataset from an international airline that sells premium-cabin seats with upgrades in the last two days before departure. The dataset contains information on both upgrades and retail sales. Analysis of the upgrade data reveals two facts: firstly, upgrades are a relevant sales channel and, secondly, the airline employs upgrades to sell premium-cabin seats at a discounted price. Examination of the retail price data shows the airline's dynamic pricing strategy across multiple products. Overall, the data align with the theory of price discrimination: the airline uses intertemporal and intratemporal price discrimination along the retail channel, while engaging in sequential price discrimination with upgrades.

2.1 Data description

The dataset contains the universe of ticket purchases for all 1,600 flights operating on two monopolistic routes between September and December 2018, for a total of 268,035 ticket purchases. While I am unable to disclose the specific routes, one route is domestic and has a duration similar to that of a flight from New York to Miami (approximately 3 hours), while the other route is international and has a duration similar to that of a flight from New York to Mexico City (approximately 6 hours). For each transaction in the dataset, there is both traveler and product information. The traveler information includes the traveler's gender, age, reason for traveling (business or leisure) and flyer-ID (whether the traveler is a member of a loyalty program). The product information includes the ticket price, timing of purchase and final travel class (economy, premium economy or business class).[4] In the case of an upgrade, the dataset records the upgrade fee. Additionally, the

[4] The dataset also contains the ticket's fare, which encompasses other features of the ticket, such as flexibility or refundability, potentially impacting ticket quality and the final price paid. However, discussions with the airline revealed that this information is not accurate, and thus, it is not utilized in this paper. For the purposes of this study, I consider all tickets within a cabin to have the same quality.

product information displays flight characteristics such as the day and time of departure and the aircraft model.

2.2 Quantity of upgrades

The airline sells a large fraction of seats by means of upgrades in the last two days before departure; this practice generates substantial revenues. I provide two pieces of evidence to support this claim. First, I present the aggregate number of transactions and revenues generated by upgrades in relation to retail sales. Second, I analyze the seat allocation over time, taking into account both retail and upgrade sales.

Compared to retail sales of premium cabin seats (premium economy and business class), upgrades serve as a relevant sales channel in terms of both transactions and revenues. According to Table 1, a large group of travelers chooses to fly in premium cabins because of the availability of the upgrade option, and the fees associated with these upgrades contribute significantly to the overall revenues of the airline. Over all premium cabin transactions, 29% of the seats in premium economy and 10% of those in business class are sold by means of an upgrade. Considering revenues, upgrading fees account for 10% of premium economy and 6% of business class revenues. In aggregate, 16% of travelers fly in premium cabins by means of upgrade purchases. The associated upgrade fees account for 6% of premium-cabin revenues and nearly 2% of total revenues.

Furthermore, when we consider sales over time, the airline allocates a large number of the seats in premium cabins in the last two days before departure. Figure 1 shows how the airline sells seats in premium cabins over time, distinguishing between retail and upgrade sales. In the last two days before departure, the number of seats allocated in premium economy through upgrades is five times the number of seats allocated by means of retail sales in the same period. We observe a similar trend for business class, where as many upgrades as retail sales take place in the last two days before departure.

2.3 Price of upgrades

The pricing data show that accessing a premium cabin is cheaper via an upgrade than via a retail purchase. I provide two pieces of evidence to support this claim. First, I analyze the time series of prices associated with accessing premium cabins. Second, I discuss the results of conditional mean regressions.

Figure 2 provides visual evidence of the discount offered through upgrades. It compares the average final price paid to access premium cabins by the time of initial ticket purchase, so that the final upgrade price is the summation of the lower-cabin ticket and the upgrade fee. Customers flying in premium economy and business class by means of an upgrade consistently enjoy discounts from the retail price.

Similarly, specification (1) examines the extent of these discounts. It compares the final prices paid to access premium cabins via retail sales and via upgrades:

$$P_{\mathrm{i}} = \alpha + \beta_{\mathrm{up}}\mathbb{I}\{\text{upgrade sale}\}_i + \epsilon_i \tag{1}$$

where P_{i} is the final price paid for ticket i to access a premium cabin. The coefficient α represents the average final price to access premium cabins through retail purchases, whereas β_{up} measures the savings (from the average final retail price) offered to travelers who access premium cabins through upgrades.

I estimate (1) separately for business and premium economy class. Table 2 shows that customers who upgrade to business class spend, on average, 13% less than they would have had they purchased a business class ticket at retail price. Similarly, Table 3 indicates that customers who upgrade to premium economy save, on average, 18% off the average retail price. Because the final paid prices might depend on the level of capacity already sold and on the time of initial purchase, selection is a concern for specification (1). Large absolute values of β_{up} might be due to variation in prices across flights with different level of sold capacity either in the last two days before departure or at the time of the initial purchase. Similarly, large absolute values of β_{up} might capture differences in prices due to variation in the time of initial purchase. I control for these effects by adding the levels of unsold capacity, time of the initial purchase and flight fixed effects as control variables to regression (1). Appendix A.2 shows that the estimation results remain robust to adding these controls.

2.4 Airline's price discrimination strategies

The evidence from retail prices, along with that related to the upgrade channel, shows that the airline's pricing strategies are consistent with its engaging in price discrimination of various degrees that have been extensively discussed in the literature. In particular, the airline engages in intertemporal, intratemporal and sequential price discrimination.

The airline allocates the majority of its seats through the retail channel. Along this channel, it implements pricing strategies consistent with well-established theories of price discrimination discussed in the existing literature. The data presented in Table 1 indicate that the airline sells 95% of its seats via the retail channel, which contributes up to 98% of total revenues. Figure 3 displays the airline's pricing patterns over time, providing support for pricing strategies that align with those described in the airline economics literature. These include the use of dynamic pricing to intertemporally price discriminate across consumers. Previous studies such as Williams (2022) and Lazarev (2013) have highlighted that customers who purchase tickets earlier in the booking horizon tend to exhibit a lower willingness to pay for travel than those making last-minute purchases. Furthermore, similarly to how airlines behave in the analysis of Aryal et al. (2023), the airline employs intratemporal price discrimination by offering seats with distinct levels of comfort at distinct prices. In this way, the airline caters to different customer segments based on their willingness to pay and value for comfort.

On the other hand, customers accessing premium cabins at a discount, after making their initial retail purchase, supports the theory of sequential price discrimination, discussed in Cui et al. (2019). In the first stage of retail sales, the airline sorts customers into different cabins based on their preferences. For instance, customers with a lower willingness to pay prefer to purchase an economy class ticket rather than a business class ticket. However, the marginal customer in economy class would have bought a business class seat had it been more affordable. Therefore, after the retail stage, the airline offers these marginal economy class customers the opportunity to upgrade to a higher cabin at a discounted price, effectively using upgrades as a sequential screening tool.

3 Model

In this section, I present a model that captures the supply and demand dynamics of the airline industry. On the supply side, I consider a capacity-constrained monopolistic airline that optimally chooses prices. In particular, in any day t of the booking horizon the airline chooses the optimal prices for two products: economy and business class. For simplicity, $t \in \{0, 1\}$, with $t = 0$ being the initial period of the booking horizon and $t = 1$ the departure day. Additionally, in $t = 1$, the airline sells upgrades, allowing those who bought an economy class ticket in $t = 0$ to pay a fee and access business class. I assume full integration of firm's pricing decisions, meaning that prices are determined by the firm's solving a dynamic programming problem that considers multiple cabins and upgrades simultaneously. This supply model reproduces relevant features of the airline's price discrimination strategies, including intertemporal, intratemporal and sequential price discrimination. On the demand side, customers randomly enter the market for airline tickets and, upon arrival, make a decision regarding whether to purchase an economy or a business class ticket or, alternatively, to not fly at all. I assume that customers do not expect the possibility of an upgrade when making their purchase decision. In $t = 1$, when the airline offers the upgrade possibility, only a fraction of customers pay attention to the offer and then consider whether to upgrade. This demand specification embeds uncertainty in the number of travelers shopping for a ticket, heterogeneity in preferences and inattention to the upgrade option.

3.1 Demand

The demand model is developed in three steps. First, following Williams (2022) and Aryal et al. (2023), I describe the retail sales market for airline tickets: how consumers arrive to the market and how they choose between buying economy or business class or not flying at all. Then, I describe how economy ticket holders decide to upgrade. Last, I describe the random coefficient structure of the demand process.

Along the retail channel, first, customers randomly arrive to the airline ticket market, and then, without expecting the upgrade option, they make their retail purchase decision. Customers arrive to the airline ticket market in each period $t \in \{0, 1\}$ of the booking horizon, starting from the initial period $t = 0$ until departure day $t = 1$. In each period t, a number N_t of customers arrive, following a Poisson process with parameter λ_t. Similarly to their counterparts in the models of Williams (2022) and Aryal et al. (2023), customers are of discrete types: they have a probability θ_t of flying for business-related reasons (H-types), and a probability of $1 - \theta_t$ of flying for leisure (L-types). Once customers enter the market, they maximize their utility by choosing between an economy or business class seat or deciding to exit the market. Customers make this decision without expecting the possibility of an upgrade. As in Aryal et al. (2023), customers differ in willingness to pay and value for comfort,[5] and goods are vertically differentiated: all customers prefer business to economy Class seats when they are sold at the same price. By defining j as the product choice, p_{EC} as the price for economy class and p_{BC} as the price for business class, I specify the utility function for individual i of type k as:

$$u_{ikj} = \begin{cases} v_{ik}\xi_{ik} - p_{BC} & \text{j= business class seat} \\ v_{ik} - p_{EC} & \text{j= economy class seat} \\ 0 & \text{j= } \emptyset \end{cases} \tag{2}$$

where the utility of not flying ($j = \emptyset$) is normalized to 0. In specification (2), v_{ik} represents the willingness to pay for traveling and $\xi_{ik} > 1$ represents the value for comfort from flying in business class. I assume that v_{ik} and ξ_{ik} are random variables independently drawn from type-specific distributions. In particular, I assume $v_{kt} \sim Exp(\beta_{kt})$ and $\xi_{kt} \sim Unif(1, \gamma_{kt})$, where β_{kt} and γ_{kt} are allowed to vary over time.

I make three assumptions when modeling the upgrade decision. First, only travelers who buy a ticket in economy class in $t = 0$ can upgrade. This assumption prevents $t = 1$ economy class buyers from upgrading to business class. Allowing only one cohort of travelers to upgrade greatly simplifies the simulation of the model, but still allows me to analyze how upgrades work. Second, travelers are naive and do not expect to be offered the upgrade option: in specification (2), utility does not include the future possibility of being upgraded. This assumption is in line with recent revenue management literature stating that, in the long run, "it is unlikely that customers will adapt strategically to upgrades"

[5] Demand shocks are a measure of the desirability of a flight. For instance, if there is a conference at the destination, business-type customers are likely to have a high WTP for flights departing before the conference; on the other hand, around Thanksgiving, leisure-type customers are willing to pay more given their need to travel over the vacation period.

(Gallego et al., 2009). Additionally, this assumption simplifies the problem since it does not require customers to form expectations about future prices. Since customers do not expect to receive the upgrade option, its only implication is a new decision problem in period $t = 1$ for those customers who had bought an economy class ticket in $t = 0$ and are attentive to the notification from the airline. Third, I assume only a fraction α_k of economy class ticket holders of type k pay attention to the notification sent by the airline. For these attentive economy class ticket holders, the decision to upgrade reduces to a comparison between the utility from flying in business class with an upgrade and the utility from flying in economy, given the retail price originally paid for an economy ticket. Therefore, customer i of type k who purchased an economy class ticket in period $t = 0$ at $p_{0,EC}$ chooses to upgrade if and only if the following inequality holds:

$$\underbrace{v_{ik}\xi_{ik} - p_{0,EC} - p_{UP}}_{\text{utility from flying in business class after an upgrade}} \geq \underbrace{v_{ik} - p_{0,EC}}_{\text{utility from flying in economy class}} \tag{3}$$

with p_{UP} being the upgrade fee. Based on expression (3), from the airline's point of view, the customer's probability of buying an upgrade, conditional on her having purchased an economy class ticket in period $t = 0$, is given by the following expression, where for ease of notation I omit indexes:

$$s_{UP}(p_{0,EC}, p_{0,BC}, p_{UP}) = \int\int \mathbb{1}\{v\xi - p_{UP} \geq v\} d\tilde{F}_{v,\xi} \tag{4}$$

with $\tilde{F}_{v,\xi}(x, y) = \mathbb{P}\left(v \leq x, \xi \leq y \middle| v - p_{0,EC} \geq \max\{v\xi - p_{0,BC}, 0\}\right)$.

In summary, the following set defines the demand for an individual flight:

$$\left\{\{\underbrace{\beta_{kt}, \gamma_{kt}}_{\Psi_t}, \alpha_k\}_{k\in\{H,L\}}, \lambda_t, \theta_t\right\}_{t\in\{0,1\}} \tag{5}$$

where $\Psi_t = \{\beta_{kt}, \gamma_{kt}\}_{k\in\{H,L\}}$ represents the random demand shock for an individual flight in period t, whose realization is ψ_t. I make parametric assumptions on the distribution of each component of the demand shock. In particular, I assume that β_{kt} is distributed according to a Normal distribution truncated at 0, $\beta_{kt} \sim TruncNorm(\mu_{kt}, \sigma_{kt})$, and γ_{kt} according to a Uniform starting from 1, $\gamma_{kt} \sim Uniform(1, \kappa_{kt})$.[6] Because of the random

[6] As highlight in Section 3.2, these parametric assumptions are known by the airline.

nature of Ψ_t, demand differs across flights; on the other hand, arrival rates, mixture of types and inattention parameters are common across flights.

3.2 Supply

I describe the airline pricing problem in three steps. First, I define the main features of the decision problem. Second, I focus on the information structure and how it changes over time. Third, I formally describe the maximization problem.

For any given flight, the airline solves a dynamic multiproduct pricing problem with capacity constraints and zero marginal costs. The problem is dynamic because the airline operates in two distinct periods: $t = 0$, the initial period, and, $t = 1$, the last period. The pricing problem is multiproduct, as the airline must determine prices for two types of products: economy and business class seats. In particular, at the beginning of any period t, the airline sets the vector p_t of prices for both economy and business class, denoted as $p_{t,EC}$ and $p_{t,BC}$, respectively. Moreover, in the last period $t = 1$, the airline sets p_{UP}, which customers who bought an economy class ticket in $t = 0$ can pay to upgrade to business class. The problem involves capacity constraints.[7] This means that the total number of economy and business class flying passengers cannot exceed the total number of seats available in these cabins, represented by $c_{0,EC}$ and $c_{0,BC}$, respectively. If, in any period t, the airline sells more tickets than the available remaining capacity ($c_{t,EC}$ and $c_{t,BC}$), it reimburses customers for their purchases of the extra tickets. I assume that the airline chooses the upgrade price simultaneously with the retail prices. However, first, the firm allocates retail passengers, and then, if there is capacity remaining in premium cabins, it sells upgrades. Last, I assume that the marginal cost of a ticket sale is zero, as in Williams (2022). This implies that costs are given by intratemporal and intertemporale shadow values of sales.

The set of state variables I_t relevant for pricing decisions in period t, which includes what the firm knows about demand, changes over time. This is summarized in the fol-

[7] Capacity constraints are exogenously given. This is in line with the models described in Williams (2022) and Aryal et al. (2022). They do not consider capacity dispatchment as a part of the airline decision problem.

lowing expression (6):

$$
\begin{aligned}
I_0 &= \{\psi_0, c_{0,EC}, c_{0,BC}\} \\
I_1 &= \{\psi_1, c_{1,BC}, p_0, \psi_0, Q^H_{0,EC}, Q^L_{0,EC}\}
\end{aligned}
\tag{6}
$$

On one hand, the airline learns about demand over time. In particular, the airline always knows the arrival rates λs, the mixtures of types θs, the levels of inattention αs and the parametric distribution of the demand shock Ψ_t. However, just before setting prices in each period t, the airline learns about the time-specific demand shock ψ_t. This means that the airline decides $t = 0$ prices based on ψ_0 and by forming expectations about Ψ_1. Then, after $t = 0$ sales take place, the airline learns about ψ_1 and takes it into account for $t = 1$ pricing decisions. On the other hand, the set of state variables I_t of the pricing problem changes over time. In particular, the set I_0 is $\{\psi_0, c_{0,EC}, c_{0,BC}\}$, which means that the airline sets initial prices based on the observed demand shock ψ_0 and the number of seats in both cabins that can be sold ($c_{0,EC}$ and $c_{0,BC}$). In contrast, the set I_1 is larger than I_0. Indeed, beyond accounting for ψ_1 and levels of unsold capacities at the beginning of $t = 1$ ($c_{1,EC}$[8] and $c_{1,BC}$), when setting $t = 1$ prices, the airline also considers initial-period prices p_0, the initial demand shock ψ_0, and the number of H-type and L-type travelers who had bought an economy class ticket in $t = 0$, $Q^H_{0,EC}, Q^L_{0,EC}$ respectively. To see why, we need to consider that in $t = 1$ the airline optimally sets prices along both the retail and the upgrade channel. When maximizing $t = 1$ expected profits and evaluating expected revenues from $t = 1$ sales, the airline takes into account that that a portion of these revenues depends on upgrade sales. Given the upgrade price, the number of travelers upgrading depends on the upgrade probability, which in turn depends on $t = 0$ prices and ψ_0, as described in expression (4), and on the total number of travelers holding an economy class ticket.

More formally, in period $t = 0$, the airline chooses $p_0 = (p_{0,EC}, p_{0,BC})$ by solving the

[8] From expression (6), the number of seats available for sale at the beginning of period *t=1*, $c_{1,EC}$, can be computed as $c_{1,EC} = c_{0,EC} - Q^H_{0,EC} - Q^L_{0,EC}$.

following dynamic problem:

$$V_0(I_0) = \max_{p_0} \mathbb{E}\Big[R_0(I_0) + \sum_{c_{1,BC}, Q^H_{0,EC}, Q^L_{0,EC}} \mathbb{P}\Big(c_{1,BC}, Q^H_{0,EC}, Q^L_{0,EC}\Big|I_0, p_0\Big)\mathbb{E}_{\Psi_1}\Big[V_1(I_1)\Big|c_{1,BC}, p_0, \psi_0, Q^H_{0,EC}, Q^L_{0,EC}\Big]\Big]. \tag{7}$$

In expression (7), I take the first expectation over the number of travelers arriving and shopping for a ticket, and the mixture of types; the term $R_0(I_0)$ represents the revenues from sales in $t = 0$ in both cabins, that cannot exceed capacity constraints; the term $\mathbb{P}(c_{1,BC}, Q^H_{0,EC}, Q^L_{0,EC}|I_0, p_0)$ is the state transition probability. The last term $\mathbb{E}_{\Psi_1}\Big[V_1(I_1)|c_{1,BC}, p_0, \psi_0, Q^H_{0,EC}, Q^L_{0,EC}\Big]$ is the expected value of $V(I_1)$, where expectations are taken with respect to Ψ_1, and $V(I_1)$ represents expected revenues at the optimal $t = 1$ prices. In particular, $V(I_1)$ is given by

$$V(I_1) = \max_{p_{1,EC}, p_{1,BC}, p_{UP}} \mathbb{E}\Big[R(I_1)\Big] \tag{8}$$

where expectations are taken with respect to the total number of customers arriving and shopping for a ticket, and the mixture of types.

As evident from expressions (8), in $t = 1$, the airline simultaneously chooses the retail and upgrade prices. However, it first allocates retail purchases and then, if there are empty premium cabin seats, allocates upgrades. This assumption reflects the airline practice of offering seats by means of upgrades and retail sales at the same time. However, by prioritizing retail sales, I do not consider the potential misallocation and consequent potential reduction in consumer surplus that may arise from sales to upgrading customers instead of higher-value retail customers.

I make three assumptions that simplify the problem on the supply side. First, the booking horizon includes two periods only. While increasing the number of periods would make the pricing problem more realistic, it would also significantly increase computational complexity. For example, let us consider a firm selling retail tickets over $t \in \{0, 1, ..., T\}$, with T being the departure day and the upgrade sales period. The firm's information set in the last period T includes previously realized prices, previously realized demand shocks and type-specific economy class sales. As the airline is forward looking, the information sets of all periods t include previously realized prices and demand shocks, along with expected future prices, demand shocks and expected future realizations of economy class retail sales. Evaluating expectations and optimal prices for any combination of variables in the state space is not feasible. Therefore, to keep the problem tractable, I consider only a two-period problem. Second, I assume that all tickets within a class (economy and business) have the same quality. In reality, tickets within the same class may have varying levels of flexibility, refundability, or additional services and thus different prices. However, by assuming uniform quality within each class,[9] I decrease the number of prices that the airline needs to determine: this reduces the dimensionality and complexity of the model. At the same time, it still allows me to analyze the main tradeoffs of the upgrade option. Last, on the demand side, I impose that customers are not strategic. That is, as discussed in Section 3, customers do not delay the purchase of a business class ticket in $t = 0$ in the hope of being offered an upgrade in $t = 1$. Additionally, customers do not time their purchase strategically. These two assumptions greatly simplify the dynamic pricing problem since I do not need to consider customers' expectations when I evaluate the firm's optimal decisions. Furthermore, the assumptions remove any commitment issues on the part of the firm. Since customers are not strategic, the firm's ability to commit to certain pricing strategies becomes irrelevant.

The mechanics of the model described in expressions (7) and (8) are similar to those in Aryal et al. (2023) and Williams (2022). In particular, Aryal et al. (2023) assume that the distribution of type-specific preferences is constant over time, and that this parameter is known to the firm form the initial period for each flight. This assumption implies that intertemporal variation in arrival rates, the composition of customer types, and level of unsold capacity jointly drive the intertemporal variation in prices within a flight. In my case, intertemporal variation in the demand shock,which then determines time–flight-specific preferences, further contributes to intertemporal price variation at the flight level. With this slightly more general modeling assumption, I can explain large within-flight intertemporal price variation, which is particularly noticeable for business class. On the other hand, this increases the complexity of my model's solution as the numerical approximation of the firm's $t = 1$ revenue expectation becomes more demanding. As in Aryal et al. (2023), Williams (2022) models type-specific preferences to be constant over time and known by the firm from the beginning of the booking horizon so that

[9] The lack of reliable data on tickets' flexibility or refundability, as described in Section 2.1, is another reason I make this assumption.

the variation in the arrival rate, mixture of types and level of unsold capacity drive the intertemporal variation in prices. Additionally, Williams (2022) includes firm-specific idiosyncratic shocks, which play a role similar to that of the demand shocks in Aryal et al. (2023) and in the model delineated in my paper. In particular, these elements explain the within-time price variation across flights with the same levels of unsold capacity.

4 Econometric model, estimation and identification

In this section, first, I review the main parametric assumptions of the structural model described in Section 3 and I discuss its identification. Second, I describe estimation.

4.1 Econometric model and identification

In this section, first, I review the main parametric assumptions of the structural model described in Section 3. Then, I discuss its identification.

According to the parametric assumptions governing the distribution of the demand shocks, the set Θ, defined as

$$\Theta = \left\{ \{\mu_{kt}, \sigma_{kt}, \kappa_{kt}, \alpha_k\}_{k \in \{H,L\}}, \lambda_t, \theta_t \right\}_{t \in \{0,1\}} \in \mathbb{R}^{18}, \tag{9}$$

fully describes the primitives of the model. In particular, the set $\{\mu_{kt}, \sigma_{kt}, \kappa_{kt}, \alpha_k\}_{k \in \{H,L\}}$ fully describes the distribution of Ψ_t.

Similarly to Aryal et al. (2023) and Williams (2022), I claim identification of the model. In particular, I use both demand-side and supply-side restrictions as identifying moments:

(i) the distribution of realized prices for $j \in \{EC, BC\}$ and $t \in \{0, 1\}$;

(ii) the average retail sales for $j \in \{EC, BC\}$, $t \in \{0, 1\}$ and $k \in \{H, L\}$;

(iii) the average fraction of business class retail sales over all sales in any period $t \in \{0, 1\}$ for any type $k \in \{H, L\}$; and

(iv) the average number of upgrades for $k \in \{H, L\}$.

The mean of the realized prices, along with their distribution, provides information about the mean, μ, and the variance, σ, of the WTP. If customers' WTP is highly volatile, optimally set prices display high variance. Similarly, if the average WTP is high, the airline charges higher prices on average. The ratio between total sales of H-types and L-types is informative of the customers mixture, θ. A large L-type sales' share indicates a large share of arriving L-type customers. When we consider a specific type, the ratio of business class retail sales to economy class sales is informative of the value for comfort, κ. Given some prices, a large proportion of business class purchases indicate a large value for comfort. With preferences identified, the average number of sales provides insight into the arrival rate, λ: if more shoppers enter the market, more sales take place. Finally, the share of upgrading customers describes the average number of attentive travelers within each type.

4.2 Estimation

In this section, I discuss how I estimate the model. In the first part, I present the simulated method of moments (SMM) estimator. In the second part, I discuss the parameter estimates and model fit.

I describe the SMM procedure in two steps: I first outline the evaluation of the empirical moments and then that of their simulated counterparts. The approach is in line with that in Aryal et al. (2023) and Nevo et al. (2016). I define the optimal estimator $\hat{\theta}$ for $\theta \in \Theta$ as the solution of the following problem:

$$\arg\min_{\theta} L(\theta) = \arg\min_{\theta} \left(m^{data} - m(\theta)\right)' \left(m^{data} - m(\theta)\right) \tag{10}$$

where $m^{data} - m(\theta)$ is a 1xM vector, which is the difference between the observed moments, m^{data}, and the corresponding simulated ones, $m(\theta)$, with M being the number of moments. In particular, I obtain $\hat{\theta}$ by minimizing $L(\theta)$ over 100,000 different values of θ with grid search. Solving problem (10) requires evaluation of both m^{data} and $m(\theta)$. I compute m^{data} from the original dataset in two steps. First, I aggregate the initial purchase dates (in terms of their distance from the departure day) into two periods: *t=0* includes all days of the booking horizon except the last two, which, in turn, constitute $t = 1$. At the same time, similarly to Aryal et al. (2023), I assume that the type k is observable from the reason for traveling.[10] Then, for all flights, I compute the total number of sales for any type and average prices in any period. The second element needed to solve problem (10) is $m(\theta)$, which requires solving the decision problem of the firm. First, for $N^s(= 100)$ different demand realizations draws from the same θ, I solve the airline's pricing problem defined by expressions (7) and (8); then, I evaluate the simulated moments $m(\theta)$. As the pricing decisions are dependent on c_{EC} and c_{BC}, I solve the pricing problem of the airline separately for various combinations of total capacities observed in the data. Last, I evaluate estimates and bootstrap standard errors by resampling flights with replacement 100 times.

Before describing the parameter estimates, I discuss three procedural assumptions employed for the estimation. First, I estimate preferences at the aircraft–route level. The estimation at the aircraft level is because of the dependence of the optimal simulated prices on the initial capacity levels. Therefore, the demand estimates reflect aircraft-specific tastes. Given the available data, estimation at the route level implies two different sets of estimates: one for the domestic and one for the international route. Moreover, estimation at the route level is consistent with practices in the industry (IATA). The second assumption is that I consider the premium economy class, when it is available, to be part of business class. Considering the firm's setting of optimal prices for two products, rather than three, simplifies the pricing problem and eases the simulation of the model. The third assumption is to impose $\mu_{H0} = \mu_{H1}$, $\sigma_{H0} = \sigma_{H1}$ and $\mu_{L0} = \mu_{L1}$. This partially reduces the flexibility of the model but simplifies minimization problem (10), as it reduces the search to one over only 15 rather than 18 parameters. This simplification still allows me to match the data generally well.

Tables 4 and 5 show estimation results for the largest aircraft in the data flying on

[10] I consider travelers reporting business to be their reason for travel to be H-types and the others L-types.

the international route.[11,12] The mean willingness to pay for an economy ticket for a flight on the international route is \$306 for H-types and \$220 for L-types. On average, there is a 40% difference in WTP between the types. As a benchmark in the literature, Aryal et al. (2023) find a 20% difference between types. However, the difference in the value for comfort per flight across types is not as large: in $t = 0$, H-types value comfort 3% more than L-types, whereas in $t = 1$, H-types value comfort 15% more than L-types. In terms of the value placed on business class compared to that on economy class, my estimates suggest that the two types value business class, on average, twice as much as economy class. In comparison to the estimates of the value for comfort in Aryal et al. (2023), where travelers value the comfort of the premium product on average 50% more than they value the regular product, my findings indicate a greater taste for quality. This is likely due to the larger average premium cabin retail prices observed in my data.

The average number of customers shopping for a ticket in *t=0* is 488, whereas in $t = 1$, it is 50. The stark decrease is consistent with how I divide the booking horizon, as $t = 0$ includes the large majority of days before departure. On the other hand, the fraction of customers traveling for business purposes slightly increases over time. This finding is consistent with findings in Aryal et al. (2023) and Williams (2022). Regarding the inattention parameters, which measure the fraction of economy class ticket holders who open the email sent by the airline containing the notification on the upgrade option, they are 0.15 and 0.09 for H-types and L-types. As a benchmark, travel industry surveys report that between 20% and 40% of people receiving a travel-related email open it.[13]

From demand estimates, I compute demand elasticities. By considering the average value of WTP and the average prices on the international route, the simulated average elasticities over time for an economy class ticket are -2.07 and -3.12 for business-type and leisure-type travelers, respectively. These estimates align with the literature. As a

[11] There are 294 flights of this kind. I estimate the model for those flights displaying $p_{EC,0} \in [200, 400]$, $p_{EC,1} \in [200, 650]$, $p_{BC,0} \in [200, 650]$ and $p_{BC,1} \in [700, 3000]$. I consider only flights displaying all these four prices, for a total number of 104 flights. I select the sample in this way in order to eliminate outliers, and reduce the size of the price grid over which the airline makes its optimal pricing decision, in order to increase the speed of the simulation.

[12] This particular aircraft has capacities of 247 for economy class, 21 for premium economy, and 31 for business class. I set the capacities to $c_{EC} = 247$ and $c_{BC} = 51$ in the simulation, as I assume that business class includes the premium economy class seats. In Appendix A.7.2, I show estimation results for the smallest aircraft in the data as well. At the beginning of Section A.7, I describe the main aircraft in the data.

[13] See: ConstantContact and campaignmonitor.com.

benchmark, Williams (2023) finds an elasticity of -3.3 for an economy class ticket, by averaging over time and across types. Aryal et al. (2023) find that price elasticity for leisure passengers is -3.9, and for business passengers, it is -0.51.

4.3 Goodness of fit: Discussion

In this section, first, I discuss the model's ability to fit the data for large aircraft flying on the international route. Second, as an illustration of robustness, I describe the relevance of the inattention parameters for fitting purposes. Finally, I describe a measure to assess the fit of the model.

Overall, the model approximates the data well for prices and sales along both the retail and upgrade channels. Figures 4, 5 and 6 show how the simulated prices and sales fit the observed data along each channel. In Figures 4 and 5, the boxplot represents the central 50% of the distribution with the median, whereas the lower and upper bounds represent the minimum and maximum. The red box represents the data, whereas the blue box represents the simulated distribution across the bootstrap estimates. Figure 6 displays the model predictions for upgrade prices and sales. The blue line represents the distribution observed in the model, whereas the red line represents the average distribution across the bootstrap estimates. The model's fitted values are generally accurate for both retail and upgrade channels, with a few exceptions. In particular, for the last period, my model predicts simulated prices that are lower than the observed ones for business class. The assumption of constant willingness to pay over time, which ensures computational tractability, likely explains this result. As a robustness check of the model's fitting ability, I estimate it for small aircraft flying on the domestic route,[14] and I show its fit in Appendix A.7.2. Despite differences in patterns in the observed data with respect to the large aircraft flying on the international route, the estimated model fits the data well, especially for retail prices and sales.

With respect to the empirical airline demand models described in Aryal et al. (2023) and Williams (2022), my framework includes inattention coefficients (α_H and α_L). These parameters play a crucial role in fitting the relevant moments in the data, in particular those related to upgrades sales. Appendix A.6 shows the importance of the inattention

[14] The most popular flight on the domestic route does not have premium economy, and it displays $c_{EC} = 153$ and $c_{BC} = 16$.

coefficients by considering the estimates and fit of a model with fully attentive customers ($\alpha_H = \alpha_L = 1$). Visually, the model with fully attentive customers fits the data worse than a model with inattentive travelers, in particular for upgrade sales. In fact, when all customers pay attention to the airline's notification related to the upgrade option, the estimated model predicts an excessively large number of upgrading customers.[15]

To formally assess the model fit, I take the same approach as Asker et al. (2014). In particular, I consider the sum of the squared differences between the observed and predicted moments, scaled by the observed moments. Evaluating the vector of observed and predicted moments, respectively x and $\hat{x}$, I compute

$$S^2 = 1 - \frac{(x-\hat{x})'(x-\hat{x})}{x'x} \tag{11}$$

as a measure of fit. In particular, the term $\frac{(x-\hat{x})'(x-\hat{x})}{x'x}$ can be interpreted as the weighted average of the square percentage differences between the observed and predicted moments. The value of S^2 is by construction less than or equal to 1, and higher values of S^2 indicate better model fit, with $S^2 = 1$ indicating a model's perfectly predicting the data. The simulation results are consistent with the visual inspection: S^2 =0.975 for the model allowing for inattention, which is larger than S^2 =0.965 for the model with fully attentive travelers.

5 Counterfactual

In this section, I describe how the upgrade option affects airline pricing decisions and welfare. In Section 5.1, I describe the role of the upgrade option as a sequential price discrimination tool and its role in managing limited inventory. In Section 5.2, firstly, I analyze the aggregate welfare consequences arising from the introduction of upgrades. Afterwards, I focus on the distributional welfare effects of upgrades between travelers and the firm.

[15] Considering inattentive customers has consequences for the counterfactual results. In particular, travelers' being fully attentive implies that the introduction of the upgrade option generates larger pricing distortions than in the case of inattentive travelers. In Appendix A.3, I explore the role of upgrades as a sequential price discrimination tool, and I compare the case where travelers are inattentive with one where travelers are fully attentive. The pricing distortions are larger with inattentive travelers.

5.1 The economics of upgrades and the role of demand shocks

To understand how upgrades work, I compare the airline's pricing decisions with and without upgrades under scenarios with and without capacity constraints. On one hand, the difference in prices between the scenarios with and the scenario without upgrades, in the absence of capacity constraints, illustrates the role of upgrades as a sequential price discrimination tool. On the other hand, comparing these differences with those observed when the airline faces capacity constraints shows how the airline uses upgrades as a way to manage inventory. In this section, I first discuss how these two channels work on average across all flights, and then I discuss how these two channels work differently conditional on the demand shock.

Table 6 shows how the upgrade option affects on average the airline's pricing problem with and without capacity constraints. In the absence of capacity constraints, there exists no connection between the retail pricing decisions in the two periods, and the upgrade option influences prices only in $t = 0$. Specifically, when setting prices in $t = 0$, the airline anticipates that some economy class ticket holders may choose to upgrade to business class if given the possibility of accessing this cabin at a reduced price. Introducing the upgrade option, then, leads to two effects. On one hand, the average opportunity cost of decreasing economy class ticket prices diminishes: lower prices for economy class incentivize more customers to purchase tickets in this cabin, thus increasing the number of potentially upgrading travelers. Introducing upgrades results in a 1.3% average decrease in economy class prices. On the other hand, the average opportunity cost of increasing business class ticket prices decreases since customers who did not purchase business class retail initially can still upgrade in $t = 1$. The introduction of upgrades leads to an average increase of 1.6% in business class prices.

Table 6 also shows how capacity constraints (with total capacity in economy class set at $c_{EC} = 247$ and in business class set at $c_{BC} = 51$) affect the airline's pricing decisions on average across flights and how they interact with the upgrade option. When the firm introduces capacity constraints, the low demand for economy class and high demand for business class relative to the constraints result in a slight change in prices for economy class and a significant increase in retail prices for business class due to scarcity. Moreover, compared to the case without constraints, the introduction of upgrades impacts both

$t = 0$ and $t = 1$ prices. Sequential price discrimination affects $t = 0$ prices, whereas the interplay between upgrades and capacity constraints affects both $t = 0$ and $t = 1$ prices. On average, as shown in Figures 7 and 8, when the airline introduces the upgrade option, it further increases—with respect to those in the scenario without capacity constraints—business class prices in $t = 0$ to reduce the number of premium-cabin sales and sellouts in the initial period and then allow for potentially higher-value retail sales in $t = 1$. In this way, the upgrade option works as an inventory management tool, leading to an average 3% increase in $t = 0$ business class prices. As a consequence, the introduction of the upgrade option reduces sales of business class in $t = 0$ more if there are constraints than if there are not. The airline compensates the loss in revenues from reduced business class sales by decreasing economy class prices in $t = 0$ slightly more (1.6%) if there are constraints, to induce more customers to access economy class and then potentially upgrade. Furthermore, when there are capacity constraints, the upgrade option also induces price changes in $t = 1$. Indeed, due to limited capacity in the premium cabin, upgrading customers compete with retail customers for the same seats in business class. This enables the airline to increase business class retail prices in $t = 1$. Moreover, due to upgrading customers leaving economy class emptier, the airline finds it optimal to decrease its $t = 1$ price to attract more passengers and fill the vacant seats.[16]

The role of demand shocks

Whether the upgrade option, beyond working as a way to sequentially price discriminate, also serves as a tool to manage inventory depends on the size of the $t = 0$ demand shocks. In case of low initial demand shocks, upgrades work mainly as a tool to sequentially price discriminate. Conversely, in case of a large demand shock, upgrades work as a tool to both sequentially price discriminate and manage inventory. In the former case, upgrades reduce the *spoilage* issue of an empty business class by allowing economy class ticket holders to fill the premium cabin, whereas in the latter, upgrades reduce the *spillage* problem of

[16] In Appendix A.8, I show the effects of the upgrade option with and without capacity constraints while considering welfare. The counterfactual results indicate that the upgrade option increases consumer and producer surplus in both scenarios.

early sellouts.[17]

When the airline faces low initial demand shocks, the probability of selling out in either cabin is negligible. This implies that capacity constraints have little effect on the pricing problem of the airline and, thus, the airline uses upgrades mainly to sequentially price discriminate among $t = 0$ travelers. As shown in Tables 7 and 8, the percentage changes in $t = 0$ prices, and then retail sales, induced by the introduction of the upgrade option are the same with and without constraints in the case of low demand shocks. Because of low demand for tickets in $t = 0$, due to the low initial demand shock, the airline faces practically no capacity constraints in $t = 1$, and therefore, upgrades have negligible effects on $t = 1$ prices. The latter fact is shown in Table 9, where the percentage changes induced by the upgrade option in $t = 1$ display high variance and are not statistically meaningful. In general, on unpopular flights, characterized by low initial demand shocks, upgrades work as a sequential price discrimination tool and mitigate spoilage issues in the premium cabin.

Conversely, in the case with a large demand shock, the airline, beyond using the upgrade option to sequentially price discriminate among $t = 0$ customers, uses it to manage limited inventory. Specifically, in the case of a large initial-period demand shock, the probability of business class selling out in $t = 0$ is large. To avoid the risk of it selling out and the airline then missing out on high-value sales in the last period, the airline exploits the upgrade option to increase business class prices in $t = 0$. With respect to the situation without capacity constraints, the increase in $t = 0$ business class prices due to the introduction of the upgrade option is larger when there are capacity constraints, as shown in Table 7. This is due to upgrades serving as a tool to manage inventory. This large increase in premium cabin prices has two effects shown in Tables 8, 10 and 11. On one hand, it reduces business class initial sales and sellouts; on the other hand, it increases economy class initial sales and sellouts because of customers buying the regular product rather than the more expensive premium one in $t = 0$. This benefits the airline as the number of potentially upgrading travelers increases. Furthermore, in $t = 1$, upgrading customers compete with retail purchasers for the same business class seats. This induces

[17] Empty flights and early sellouts are widely recognized problems in the airline industry: “Hold inventory (high) for too long, and they could risk having a plane depart with empty seats (spoilage). The stakes are incredibly high—sell too much too early at a lower price, and airlines might sell out too early missing out on high yielding last-minute sales (spillage)” (Alaska Airlines).

the airline to increase $t = 1$ business class retail prices, as shown in Table 9. In case of a large demand shock, many economy class ticket holders upgrade: with respect to the scenario without upgrades; this leads to more business class sellouts and fewer economy class sellouts at the end of $t = 1$, as shown in Table 11. In general, on popular flights, characterized by large demand shocks, upgrades, beyond serving as a sequential price discrimination tool, work as an inventory management tool and mitigate spillage issues in the premium cabin.

5.2 Welfare effects of the upgrade option

In this section, first, I describe how the introduction of the upgrade option affects total welfare; then I discuss how it affects travelers and the airline separately. To provide further context, in Appendix A.4, I present a comparison of the welfare gains from the upgrade option with the gains from dynamic pricing – with respect to uniform pricing – and a free upgrade policy.

Aggregate welfare effects. The upgrade option increases the surplus of both travelers and the airline, thereby increasing efficiency, as shown in Table 12. In particular, consumer surplus increases by an average of 1.5% per flight, proving that the welfare gains enjoyed by upgrading customers outweigh the consumer welfare losses arising from higher business class prices. By looking at the average number of passengers boarding the plane when the airline does not implement upgrades, such an increase in consumer surplus is equivalent to a \$6.5 subsidy to all boarding travelers. On the firm side, the airline's surplus increases by 2% per flight, primarily because of the substantial revenues generated by upgrade fees. As a benchmark, the global profit margin in the airline industry in 2023 averages around 1.2% (IATA). Moreover, Cui et al. (2019) find a 4% increase in revenues when an airline introduced add-on products[18].

Effects on consumers. The introduction of the upgrade option modifies the airline's pricing problem, and thus it affects how the same customer behaves in the scenarios with

[18] Cui et al. (2019) study an airline that allowed economy class ticket holders to upgrade to premium economy. There is no retail channel for premium economy. Their framework misses the interaction between capacity constraints and the upgrade option for the premium product.

and without upgrades and how welfare is distributed among travelers. Tables 13 and 14 provide a summary of the changes in customers' decisions based on simulations of 500 flights. Table 15 illustrates how consumer welfare and its distribution across cabins change when the upgrade option is removed.

Table 13 shows how the upgrade option affects the allocation of passengers across cabins. For instance, in the presence of the upgrade option, 181,032 customers enter the market for airline tickets but do not purchase any tickets. However, when the upgrade option is eliminated and different prices come into play, 861 of these customers switch to economy class, and 52 switch to business class. Table 13 aligns with Table 12, showing that the upgrade option reduces business class retail sales and increases economy class retail sales. Despite this, the number of passengers flying in business class is higher with the upgrade option because of the 3,568 upgrading travelers. Furthermore, the upgrade option allows more travelers to fly. Table 14 presents the changes in consumer behavior in percentage terms. It indicates the percentage of customers who change their decisions when the upgrade option is eliminated. In particular, 99.5% of customers who do not buy a ticket when there are upgrades continue to prefer the outside option even without upgrades. However, 0.5% of them switch to economy class in the absence of the upgrade option. According to Tables 13 and 14, 2.5% of customers (88 in total) who upgrade to business class switch to a business class retail purchase when the upgrade option is not available. This demonstrates the "cannibalization effect" induced by upgrades, as the introduction of the upgrade option eliminates part of the business class retail sales.

Table 15 illustrates how consumer welfare and its distribution across cabins change when the upgrade option is removed. The simulation results indicate that, on average, customers gain $1,227 (+1.5%) per flight from the upgrade option. The upgrade option primarily benefits those customers in economy class with a relatively high WTP and value for comfort, who are able to upgrade to business class. When the upgrade option is eliminated, two distinct segments of customers emerge from this group. The first segment consists of customers with a relatively lower WTP and value for comfort, who end up flying in economy class. The second segment comprises customers with a relatively higher WTP and value for comfort, who immediately purchase business class tickets at retail prices. Both segments enjoy the benefits of the upgrade option, as they gain access to the premium cabin at a discounted price. Furthermore, the upgrade option benefits those customers with a low WTP and low value for comfort, who choose to purchase economy class tickets under both scenarios. These customers benefit from the upgrade option as it leads to lower prices in economy class. On the other hand, the upgrade option reduces the welfare of customers who choose to buy business class via the retail channel in both scenarios. These customers, with a very high WTP and value for comfort, experience a loss in welfare due to the higher business class prices when the upgrade option is in place.

Effects on the firm. The upgrade option modifies how the airline generates revenues between the two cabins, as shown in Table 16. Introducing upgrades increases revenues from the retail sales for economy class but starkly decrease revenues from retail sales for business class. These effects are attributable to the lower prices in economy class, resulting in higher sales, and to the higher prices in business class, leading to cannibalization of retail sales. However, the net effect of the upgrade option on revenues is positive. Table 17 presents the changes in the distribution of revenues across products resulting from the introduction of the upgrade option. Similarly to the original data in Table 1, up to 2.4% of total revenues are derived from upgrade fees.

6 Conclusion

This paper investigates the welfare implications of introducing upgrades within the airline industry. To achieve this, I analyze proprietary data from an international airline that employs upgrades to allocate premium-cabin seats. The data show how the airline uses upgrades and allow me to estimate a structural model. The empirical analysis shows that upgrades are a relevant sales channel for premium products and that the airline employs them for price discrimination and inventory management purposes. After estimating a structural model that captures key aspects of airline pricing decisions, including multiproduct offering, dynamic pricing, and capacity constraints, I quantify the effect of upgrades on welfare through counterfactual simulations. My results indicate that, on average, both consumers and the firm benefit from the upgrade option.

There are many interesting avenues for future research based on the findings presented in this paper. One natural direction is to explore the interaction between upgrade mechanisms and competition. While my analysis focuses on a monopolistic seller implementing upgrades, it would be valuable to examine the effect of upgrade mechanisms in competitive settings. Upgrades might soften competition among firms and increase market power, with potentially negative consequences for customers. Additionally, considering the impact of strategic customer behavior would offer insights. Strategic customers may time their initial purchase decisions[19] or delay the purchase of premium products to take advantage of future upgrade discounts. Although these strategic dimensions are not included in the current analysis for computational tractability, they might influence how the upgrade option affects welfare.

[19] As in Lazarev (2013).

References

[1] aa.com. Upgrade with miles. https://www.aa.com/i18n/aadvantage-program/miles/redeem/award-travel/upgrade-with-miles.jsp, 2023. Accessed on 2023-10-24.

[2] Inaki Aguirre, Simon Cowan, and John Vickers. Monopoly price discrimination and demand curvature. *American Economic Review*, 100(4):1601–1615, 2010.

[3] allianztravelinsurance.com. 5 upsells to avoid when renting a car. https://www.allianztravelinsurance.com/travel/rental-cars/rental-car-upsells.htm#:~:text=Car%20rental%20companies%20make%2010,or%20a%20hard-sell.%22. Accessed on 2023-10-24.

[4] Gaurab Aryal, Charles Murry, and Jonathan W Williams. Price discrimination in international airline markets. *The Review of Economic Studies*, page rdad037, 03 2023.

[5] John Asker, Allan Collard-Wexler, and Jan De Loecker. Dynamic inputs and resource (mis) allocation. *Journal of Political Economy*, 122(5):1013–1063, 2014.

[6] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–957, 2015.

[7] campaignmonitor.com. Ultimate email marketing benchmarks for 2022: By industry and day. https://www.campaignmonitor.com/resources/guides/email-marketing-benchmarks/, 2022. Accessed on 2023-10-24.

[8] Ambarish Chandra. Price discrimination along multiple dimensions: New evidence from a regional airline. *Available at SSRN 3694848*, 2020.

[9] Sungjin Cho, Gong Lee, John Rust, and Mengkai Yu. Optimal dynamic hotel pricing. In *2018 Meeting Papers*, volume 179, 2018.

[10] Gregory S. Crawford and Matthew Shum. Monopoly quality degradation and regulation in cable television. *The Journal of Law and Economics*, 50(1):181–219, 2007.

[11] Ruomeng Cui and Hyoduk Shin. Sharing aggregate inventory information with customers: Strategic cross-selling and shortage reduction. *Management Science*, 64(1):381–400, 2018.

[12] Yao Cui, Izak Duenyas, and Ozge Sahin. Pricing of conditional upgrades in the presence of strategic consumers. *Management Science*, 64(7):3208–3226, 2018.

[13] Yao Cui, A Yeşim Orhun, and Izak Duenyas. How price dispersion changes when upgrades are introduced: Theory and empirical evidence from the airline industry. *Management Science*, 65(8):3835–3852, 2019.

[14] James D Dana Jr and Kevin R Williams. Intertemporal price discrimination in sequential quantity-price games. *Marketing Science*, 41(5):966–981, 2022.

[15] delta.com. All the ways to upgrade. https://www.delta.com/us/en/skymiles/program-resources/ways-to-upgrade, 2023. Accessed on 2023-10-24.

[16] Xavier Gabaix. Behavioral inattention. In *Handbook of behavioral economics: Applications and foundations 1*, volume 2, pages 261–343. Elsevier, 2019.

[17] Guillermo Gallego and Catalina Stefanescu. Upgrades, upsells and pricing in revenue management. *Available at SSRN 1334341*, 2009.

[18] IATA.org. Air travel demand. https://www.iata.org/en/iata-repository/publications/economic-reports/air-travel-demand/, 2008. Accessed on 2023-10-24.

[19] IATA.org. The importance of air transport to the united states. https://www.\.org/en/iata-repository/publications/economic-reports/the-united-states--value-of-aviation/#:~:text=The%20air%20transport%20industry%2C%20including,totaling%20to%20US%20%24779%20billion., 2018. Accessed on 2023-10-24.

[20] IATA.org. Airline profitability outlook strengthens. https://www.iata.org/en/pressroom/2023-releases/2023-06-05-01/#:~:text=Achieving%20profitability%20at%20an%20industry,a%201.2%25%20net%20profit%20margin., 2023. Accessed on 2023-10-24.

[21] IATA.org. Airline profitability outlook strengthens. https://www.iata.org/en/pressroom/2023-releases/2023-06-05-01/, 2023. Accessed on 2023-10-24.

[22] insideflyer.co.uk. insideflyer. https://insideflyer.co.uk/2017/04/one-stop-guide-upgrading-qatar-airways/. Accessed on 2023-10-24.

[23] klm.com. Upgrade options. https://www.klm.com/information/manage-booking/upgrades, 2023. Accessed on 2023-10-24.

[24] knowledgebase.constantcontact.com. Ultimate email marketing benchmarks for 2022: By industry and day. https://knowledgebase.constantcontact.com/articles/KnowledgeBase/5409-average-industry-rates?lang=en_US, 2022. Accessed on 2023-10-24.

[25] John Lazarev. The welfare effects of intertemporal price discrimination: an empirical analysis of airline pricing in us monopoly markets. *New York University*, 2013.

[26] Phillip Leslie. Price discrimination in broadway theater. *RAND Journal of Economics*, pages 520–541, 2004.

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

[28] lufthansa.com. Upgrade options. https://www.lufthansa.com/cy/en/upgrade-options, 2023. Accessed on 2023-10-24.

[29] macrotrends. macrotrends.net. https://www.macrotrends.net/. Accessed on 2023-10-24.

[30] McKinsey.com. Leading from the front line: How airlines can boost ancillary revenues. https://www.mckinsey.com/industries/travel-logistics-and-infrastructure/our-insights/leading-from-the-front-line-how-airlines-can-boost-ancillary-revenues, 2019. Accessed on 2023-10-24.

[31] Aviv Nevo, John L. Turner, and Jonathan W. Williams. Usage-based pricing and demand for residential broadband. *Econometrica*, 84(2):411–443, 2016.

[32] OAG.com. Shaping airline retail: The unstoppable rise of ancillaries. https://www.oag.com/blog/shaping-airline-retail-unstoppable-rise-ancillaries. Accessed on 2023-10-24.

[33] skypearl.csair.com. China southern skypearl award upgrade. https://skypearl.csair.com/kr/redemption/upgrade_cz.html, 2023. Accessed on 2023-10-24.

[34] tripadvisor.co.uk. received email for upgrade- is it worth it? https://www.tripadvisor.co.uk/ShowTopic-g1-i10702-k8061779-Received_email_for_upgrade_is_it_worth_it-Air_Travel.html, 2015. Accessed on 2023-10-24.

[35] united.com. All about upgrades. https://www.united.com/ual/en/us/fly/mileageplus/awards/request-an-upgrade.html, 2023. Accessed on 2023-10-24.

[36] Hal Varian. Price discrimination and social welfare. *American Economic Review*, 75(4):870–75, 1985.

[37] volantio.com. Alaska airlines. https://www.volantio.com/alaska-airlines, 2013-2020. Accessed on 2023-10-24.

[38] Kevin R Williams. The welfare effects of dynamic pricing: Evidence from airline markets. *Econometrica*, 90(2):831–858, 2022.

[39] Övünç Yılmaz, Mark Ferguson, Pelin Pekgün, and Guangzhi Shang. Strategic behavior for hotel standby upgrade programs: Empirical evidence and pricing implications. *Journal of Operations Management*, 68(6-7):675–701, 2022.

Tables

Table 1: Distribution of transactions and revenues across products and sales channels

	Economy	Premium Economy	Business
Transactions	239,771	7,820	21,278
Retail	239,771 (100%)	5,542 (71%)	19,030 (90%)
Upgrades		2,278 (29%)	2,248 (10%)
Revenues (000)	66,490$	4,653$	23,437$
Retail (000)	66,490$ (100%)	3,841$ (90%)	22,365$ (94%)
Upgrades (000)		474$ (10%)	1,298$ (6%)

Notes: Revenues are expressed in thousand $. Percentages are with respect to total transactions (or revenues) of the corresponding class. When considering upgrades, I include also auction upgrades.

Table 2: Business class, evidence of discount

Variable	Estimate
Average retail price (no upgrade)	1,337.98***
	(8.82)
Savings due to upgrades	-171.98***
	(13.46)
Number of transactions	17,413

Notes: Results are in $ and computed using sales of business class seats. When considering the final price paid after upgrades, I also include auction upgrades. I drop sales to customers with flyers-IDs. Standard errors are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors are clustered at the flight level.

Table 3: Premium economy class, evidence of discount

Variable	Estimate
Average retail price (no upgrades)	701.78***
	(6.88)
Savings due to Upgrade	-128.93***
	(4.55)
Number of transactions	7,455

Notes: Results are in $ and computed using sales from premium economy class seats. When considering the final price paid after upgrades, I also include auction upgrades. I drop sales to customers with flyers-IDs. Standard errors are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard Errors are clustered at the flight level.

Table 4: Preferences – large aircraft on the international route

t=0		t=1	
Parameter	Estimate	Parameter	*Estimate*
(μ_{H0}, σ_{H0})	(306.8,95.6)	(μ_{H1}, σ_{H1})	(306.8,95.6)
	((1.99),(4.97))		((1.99),(4.97))
κ_{H0}	3.3	κ_{H1}	3.7
	(0.0)		(0.0)
(μ_{L0}, σ_{L0})	(220.4,230.0)	(μ_{L1}, σ_{L1})	(220.4,300.0)
	((21.86),(0.0))		((21.86),(0.0))
κ_{L0}	3.22	κ_{L1}	3.2
	(0.01)		(0.02)

Notes: Bootstrap estimates and standard errors for flights in the international route on large aircraft (with $c_{EC} = 247$ and $c_{BC} = 51$). The data used for estimation include 98 flights. I use 100 bootstrap samples, by resampling at the flight-level.

Table 5: Arrival process and inattention – large aircraft on the international route

Parameter	Estimate
λ_0	488
	(0.0)
λ_1	50
	(0.0)
θ_0	0.1
	(0.0)
θ_1	0.2
	(0.01)
α_H	0.15
	(0.01)
α_L	0.09
	(0.01)

Notes: Bootstrap estimates and standard errors for flights on the international route on large aircraft (with $c_{EC} = 247$ and $c_{BC} = 51$). The data used for estimation include 98 flights. I use 100 bootstrap samples by resampling at the flight level.

Table 6: Counterfactual in levels, role of capacity constraints

	WITH capacity constraints		WITHOUT capacity constraints	
	With upgrades (1)	Without upgrades (2)	With upgrades (3)	Without upgrades (4)
$p_{EC,0}$	303	308	313	317
	(3.33)	(2.95)	(2.95)	(2.49)
$p_{EC,1}$	418	422	430	430
	(5.29)	(5.64)	(5.28)	(5.28)
$p_{BC,0}$	1,305	1,268	1,152	1,134
	(12.77)	(11.32)	(14.23)	(11.96)
$p_{BC,1}$	996	983	848	848
	(13.92)	(24.15)	(15.61)	(15.61)
p_{UP}	297		268	
	(134)		(115)	
$passengers_{EC}$	164.5	167.2	145.69	150.27
	(2.85)	(3.24)	(2.25)	(2.82)
$passengers_{BC}$	29.3	23.82	42.36	35.13
	(0.86)	(0.83)	(1.73)	(1.61)
$upgrades$	7.11		7.78	
	(0.19)		(0.25)	

Notes: I simulate the estimates and bootstrap standard errors for large aircraft on the international route. I use 10 bootstrap samples, each simulating 500 aircraft. The scenario with capacity constraints considers $c_{EC} = 247$ and in business class $c_{BC} = 51$. The variable $passengers_k$ indicates the number of passengers flying in cabin k.

Table 7: Price change in $t = 0$, sequential price discrimination and inventory management

Demand shock	WITH capacity constraints – Overall change		WITHOUT capacity constraints – Change due to SPD	
	$\%\Delta p_{0,EC}$	$\%\Delta p_{0,BC}$	$\%\Delta p_{0,EC}$	$\%\Delta p_{0,BC}$
H,H	-0.002	0.042	0.005	0.023
	(0.012)	(0.015)	(0.006)	(0.011)
H,M	0.004	0.035	-0.003	0.001
	(0.007)	(0.008)	(0.002)	(0.003)
M,H	-0.032	0.034	-0.034	0.000
	(0.015)	(0.008)	(0.015)	(0.001)
H,L	0.000	0.005	0.001	0.002
	(0.000)	(0.008)	(0.002)	(0.003)
L,H	-0.106	0.084	-0.106	0.084
	(0.029)	(0.041)	(0.002)	(0.041)
M,M	-0.009	0.031	-0.010	0.013
	(0.003)	(0.008)	(0.004)	(0.006)
M,L	-0.005	0.038	-0.011	0.036
	(0.005)	(0.016)	(0.005)	(0.018)
L,M	-0.005	0.028	-0.005	0.028
	(0.011)	(0.013)	(0.011)	(0.013)
L,L	0.016	0.012	0.016	0.012
	(0.017)	(0.012)	(0.017)	(0.012)

Notes: I simulate 500 flights operating on the international route, and I report the percentage change induced by the introduction of the upgrade option. The overall change considers the percentage difference in $t = 0$ prices arising from the introduction of the upgrade option when the airline faces capacity constraints, thus $\Delta p = \frac{p^u - p^{NOu}}{p^{NOu}}$, where p^u is the price when the upgrade option is available, whereas p^{NOu} is the price when the upgrade option is not available. Columns under "change due to SPD" indicate the percentage change due to the introduction of the upgrade option when the airline does not face capacity constraints and, thus, associated with the use of upgrades as a sequential price discrimination (SPD) tool. Demand shocks are the $t = 0$ leisure-type demand shocks, in the form of $\beta_{leisure,0}, \gamma_{leisure,0}$, where H represents a realization of the random coefficient in the top 20% of its distribution, L a realization in the bottom 20%, and M a realization in the rest. Bootstrapped standard errors over 10 samples are recorded.

Table 8: Quantity change in $t = 0$, sequential price discrimination and inventory management

Demand shock	WITH capacity constraints – Overall change		WITHOUT capacity constraints – Change due to SPD	
	$\%\Delta q_{0,EC}$	$\%\Delta q_{0,BC}$	$\%\Delta q_{0,EC}$	$\%\Delta q_{0,BC}$
H,H	0.027	-0.067	0.021	-0.027
	(0.014)	(0.026)	(0.008)	(0.012)
H,M	0.015	-0.060	0.007	-0.004
	(0.005)	(0.012)	(0.007)	(0.005)
M,H	0.110	-0.088	0.093	-0.031
	(0.030)	(0.007)	(0.039)	(0.017)
H,L	0.000	-0.002	-0.000	-0.002
	(0.000)	(0.007)	(0.001)	(0.005)
L,H	0.445	-0.191	0.435	-0.186
	(0.111)	(0.066)	(0.106)	(0.063)
M,M	0.035	-0.060	0.025	-0.025
	(0.111)	(0.066)	(0.011)	(0.011)
M,L	0.015	-0.056	0.023	-0.059
	(0.010)	(0.020)	(0.010)	(0.028)
L,M	0.033	-0.047	0.032	-0.046
	(0.031)	(0.018)	(0.031)	(0.018)
L,L	-0.017	-0.009	-0.017	-0.009
	(0.026)	(0.015)	(0.026)	(0.015)

Notes: I simulate 500 flights operating on the international route, and I report the percentage change induced by the introduction of the upgrade option. The overall change considers the percentage difference in $t = 0$ quantities arising from the introduction of the upgrade option when the airline faces capacity constraints; thus, $\Delta q = \frac{q^u - q^{NOu}}{q^{NOu}}$, where q^u is the realized quantity when the upgrade option is available, whereas q^{NOu} is the quantity when the upgrade option is not available. Columns under "change due to SPD" indicate the percentage change due to the introduction of the upgrade option when the airline does not face capacity constraints and, thus, associated to the use of upgrades as a sequential price discrimination (SPD) tool. Demand shocks are the $t = 0$ leisure-type demand shocks, in the form of $\beta_{leisure,0}, \gamma_{leisure,0}$, where H represents a realization of the random coefficient in the top 20% of its distribution, L a realization in the bottom 20% and M a realization in the rest. Bootstrapped standard errors over 10 samples are recorded.

Table 9: Leisure-type demand shock and the effect of introducing upgrades on $t = 1$ prices

Demand shock	$\%\Delta p_{1,EC}$
H,H	0.016
	(0.061)
H,M	0.016
	(0.016)
M,H	0.011
	(0.025)
H,L	0.007
	(0.033)
L,H	-0.008
	(0.052)
M,M	0.004
	(0.012)
M,L	0.023
	(0.019)
L,M	-0.007
	(0.015)
L,L	0.005
	(0.046)

Demand shock	$\%\Delta\ p_{1,BC}$
H,H	0.145
	(0.138)
H,M	0.114
	(0.060)
M,H	0.096
	(0.041)
H,L	0.015
	(0.042)
L,H	0.049
	(0.063)
M,M	0.064
	(0.024)
M,L	0.025
	(0.030)
L,M	0.035
	(0.036)
L,L	0.061
	(0.069)

Notes: I simulate 500 flights operating on the international route under the two scenarios: with and without upgrades. In both cases, the airline faces the same demand shocks and capacity constraints. I evaluate optimal prices for any flight under the two scenarios and then consider their relative difference, in particular the *Change in* $p = \frac{p^U - p^{NoU}}{p^{NoU}}$, with p^U being the scenario with upgrades and p^{NoU} being the scenario without upgrades. Demand shocks are the $t = 0$ leisure-type demand shocks, in the form of $\beta_{leisure,0}, \gamma_{leisure,0}$, where H represents a realization of the random coefficient in the top 20% of its distribution, L a realization in the bottom 20% and M a realization in the rest. Bootstrapped standard errors over 10 samples are recorded.

Table 10: Leisure-type demand shock and the effect of introducing upgrades on $t = 1$ retail sales

Demand Shock	$\%\Delta q_{1,EC}$
H,H	0.210
	(0.123)
H,M	0.119
	(0.073)
M,H	0.209
	(0.131)
H,L	0.080
	(0.109)
L,H	0.192
	(0.136)
M,M	0.225
	(0.095)
M,L	0.153
	(0.074)
L,M	0.176
	(0.084)
L,L	0.096
	(0.077)

Demand Shock	$\%\Delta q_{1,BC}$
H,H	-0.175
	(0.095)
H,M	-0.071
	(0.058)
M,H	-0.046
	(0.088)
H,L	-0.008
	(0.030)
L,H	0.011
	(0.065)
M,M	-0.029
	(0.046)
M,L	0.016
	(0.020)
L,M	0.003
	(0.028)
L,L	-0.022
	(0.026)

Notes: I simulate 500 flights operating on the international route under two scenarios: with and without upgrades. In both cases, the airline faces the same demand shocks. Demand shocks are the $t = 0$ leisure-type demand shocks, in the form of $\beta_{leisure,0}, \gamma_{leisure,0}$, where H represents a realization of the random coefficient in the top 20% of its distribution, L a realization in the bottom 20% and M a realization in the rest. I evaluate retail sales for any flight under the two scenarios and then consider their difference, in particular, *Change in q* $= \frac{q^U - q^{NoU}}{q^{NoU}}$, with q^U being the quantity in the scenario with upgrades and q^{NoU} the quantity in the scenario without upgrades. Bootstrapped standard errors over 10 samples are recorded.

Table 11: Leisure-type demand shock and the effect of introducing upgrades on sellouts

Demand Shock	$\%\Delta sellout_{0,EC}$
H,H	0.02870
	(0.05563)
H,M	0.01446
	(0.00965)
M,H	0.00000
	(0.00000)
H,L	0.00000
	(0.00000)
L,H	0.00000
	(0.00000)
M,M	0.00000
	(0.00000)
M,L	0.00378
	(0.00613)
L,M	0.00000
	(0.00000)
L,L	0.00000
	(0.00000)

Demand Shock	$\%\Delta sellout_{0,BC}$
H,H	-0.17494
	(0.06752)
H,M	-0.07696
	(0.03174)
M,H	-0.09009
	(0.03779)
H,L	0.00000
	(0.00000)
L,H	0.00000
	(0.00000)
M,M	-0.01770
	(0.00620)
M,L	0.00000
	(0.00000)
L,M	0.00000
	(0.00000)
L,L	0.00000
	(0.00000)

Demand Shock	$\%\Delta sellout_{1,EC}$
H,H	-0.00305
	(0.05567)
H,M	-0.04342
	(0.03066)
M,H	0.00000
	(0.00000)
H,L	-0.30982
	(0.07834)
L,H	0.00000
	(0.00000)
M,M	-0.00394
	(0.00398)
M,L	-0.01508
	(0.01322)
L,M	0.00000
	(0.00000)
L,L	0.00000
	(0.00000)

Demand Shock	$\%\Delta sellout_{1,BC}$
H,H	0.25012
	(0.13784)
H,M	0.18193
	(0.04648)
M,H	0.16662
	(0.05424)
H,L	0.00417
	(0.01318)
L,H	0.00000
	(0.00000)
M,M	0.05332
	(0.01599)
M,L	0.00000
	(0.00000)
L,M	0.00000
	(0.00000)
L,L	0.00000
	(0.00000)

Notes: I simulate 500 flights operating on the international route under two scenarios: with and without upgrades. In both cases, the airline faces the same demand shocks. Demand shocks are the $t = 0$ leisure-type demand shocks, in the form of $\beta_{leisure,0}, \gamma_{leisure,0}$, where H represents a realization of the random coefficient in the top 20% of its distribution, L a realization in the bottom 20% and M a realization in the rest. I evaluate the fraction (over the 500 simulations) of flights that sell out; in particular, *Change in* $sellout_k = sellout_k^U - sellout_k^{NoU}$, with $sellout_k^U$ being the fraction of flights that sell out in cabin k in the scenario with the upgrade option and $sellout_k^{NoU}$ being the fraction of flights that sell out in cabin k in the scenario without the upgrade option. Bootstrapped standard errors over 10 samples are recorded.

Table 12: Counterfactual in levels, introducing the upgrade option

	With upgrades	Without upgrades	$\Delta = With - Without$ upgrades
$p_{EC,0}$	303	308	-5
	(3.33)	(2.95)	(1.05)
$p_{EC,1}$	418	422	-4
	(5.29)	(5.64)	(2.57)
$p_{BC,0}$	1,305	1,268	37
	(12.77)	(11.32)	(3.94)
$p_{BC,1}$	996	983	13
	(13.92)	(24.15)	(15.74)
p_{UP}	297		
	(134)		
$passengers_{EC}$	164.5	167.2	-2.698
	(2.85)	(3.24)	(0.51)
$passengers_{BC}$	29.3	23.82	5.478
	(0.86)	(0.83)	(0.21)
upgrades	7.11		
	(0.19)		
$sellout_{EC}$	0.01	0.03	-0.02
	(0.004)	(0.009)	(0.007)
$sellout_{BC}$	0.166	0.096	0.07
	(0.016)	(0.015)	(0.01)
CS	81,978	80,751	1,227
	(3,017)	(3,049)	(178)
PS	80,451	78,897	1,553
	(2,251)	(2,214)	(146)
TS	162,430	159,649	2,780
	(5,244)	(5,239)	(280)

Notes: I evaluate the estimates and bootstrap standard errors for large aircraft on the international route. I use 10 bootstrap samples, each simulating 500 aircraft. The variable $passengers_k$ indicates the number of passengers flying in cabin k. The variable $sellouts_k$ indicates the fraction of flights (over the 500 simulated flights) that sell out in cabin k. **CS**, **PS** and **TS** indicate average (per flight) consumer, producer and total surplus, respectively.

Table 13: Counterfactual, change in the absolute number of retail sales

	Eliminating upgrades...	**Without Upgrades** Outside Option	*EC*	*BC*	*total*
With Upgrades	Outside Option	181,032	861	52	181,945
		(1,588)	(83)	(11)	(1,593)
	EC	2,050	79,200	1,013	82,263
		(315)	(1,738)	(48)	(1513)
	$EC + UP$	57	3,423	88	3,568
		(15)	(83)	(13)	(92)
	BC	132	192	10,725	11,049
		(28)	(30)	(401)	(395)
	total	183,271	83,676	11,878	
		(1,776)	(1,738)	(408)	

Notes: I simulate 500 flights under two scenarios: with and without upgrades for 10 bootstrap samples. I report bootstrap estimates and standard errors. I evaluate how the behavior of consumers changes when the upgrade option is eliminated.

Table 14: Counterfactual, change in retail sales as percentage of original purchases

	Eliminating upgrades...	**Without Upgrades** Outside Option	*EC*	*BC*
With Upgrades	Outside Option	99.5%	0.5%	0.0%
		(0.0)	(0.0)	(0.0)
	EC	2.5%	96.3%	1.2%
		(0.4)	(0.4)	(0.1)
	$EC + UP$	1.6%	95.9%	2.5%
		(0.4)	(0.5)	(0.4)
	BC	1.2%	1.7%	97.1%
		(0.3)	(0.3)	(0.5)

Notes: I simulate 500 flights under two scenarios: with and without upgrades for 10 bootstrap samples. I report bootstrap estimates and standard errors. I evaluate how the behavior of consumers changes in percentage terms when the upgrade option is eliminated.

Table 15: Counterfactual, consumer surplus net effects

Eliminating Upgrades $\nearrow$	Outside Option	EC	BC
Outside Option	0	-39,666	-5,386
	(0)	(8,063)	(1,983)
EC	110,666	153,748	-74,069
	(13,496)	(41,835)	(9,154)
$EC+UP$	8,132	531,122	36,301
	(2,413)	(21,090)	(5,938)
BC	174,542	15,366	-297,130
	(37,143)	(3,996)	(22,976)

Notes: I simulate 500 flights under the scenarios with and without upgrades for 10 bootstrap samples. I report bootstrap estimates and standard errors. I evaluate consumer surplus for any passenger in all 500 flights under the two scenarios and then consider their difference: $CS^{\text{scenario with upgrades}} - CS^{\text{scenario without upgrades}}$. Results are in $.

Table 16: Producer surplus counterfactual, aggregate revenues across cabins

Cabin	With Upgrades	Without upgrades	Δ
EC	53.588$	52.759$	+829$
	(1.299)	(1.304)	(89)
BC	24.929$	26.139$	-1.210$
	(1.067)	(1.099)	(120)
UP	1.935$	0$	+1.935$
	(71)	0	
total	80.452$	78.898$	+1.553$
	(2.136)	(2.101)	(138.9)

Notes: I simulate 500 flights under the scenarios: with and without upgrades for 10 bootstrap samples. I report bootstrap estimates and standard errors. I evaluate producer surplus for any flight under the scenarios with and without upgrades; then, I consider the average across the 500 flights. Results are in $.

Table 17: Producer surplus counterfactual, distribution of revenues across cabins

Cabin	With Upgrades	Without upgrades
EC	0.666	0.669
	(0.007)	(0.008)
BC	0.310	0.331
	(0.007)	(0.008)
UP	0.024	0
	(0.001)	(0)

Notes: I simulate 500 flights under the scenarios: with and without upgrades for 10 bootstrap samples. I report bootstrap estimates and standard errors. I evaluate the average distribution of producer surplus across products per scenario.

Figures

Figure 1: Distribution of sales over time

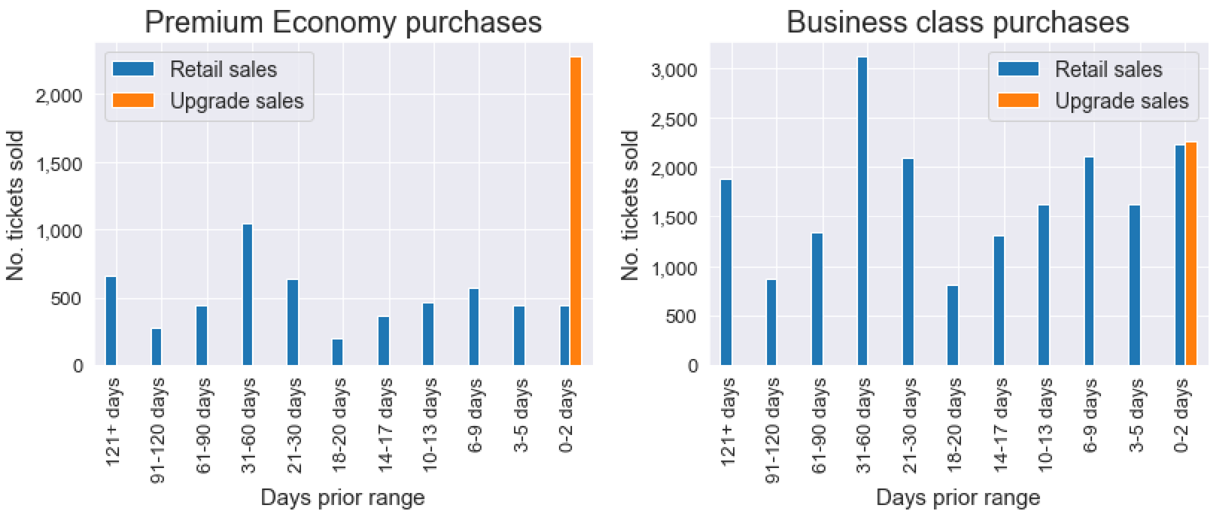

Notes: The horizontal axis shows the booking horizon by period before departure. The vertical axis displays the total number of tickets purchased over all flights in the dataset. I exclude travelers belonging to frequent flyer programs.

Figure 2: Discount over time

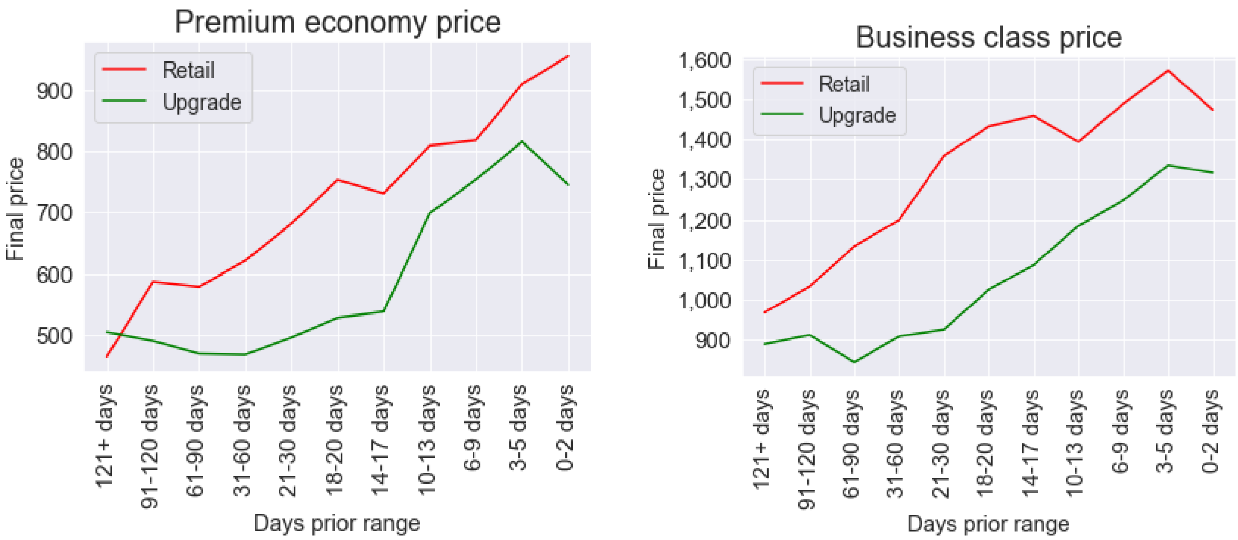

Notes: The horizontal axis shows the booking horizon split by period before departure. The vertical axis displays the final paid price paid to access premium cabins. I exclude travelers belonging to frequent flyer programs.

Figure 3: Evolution of retail prices over time

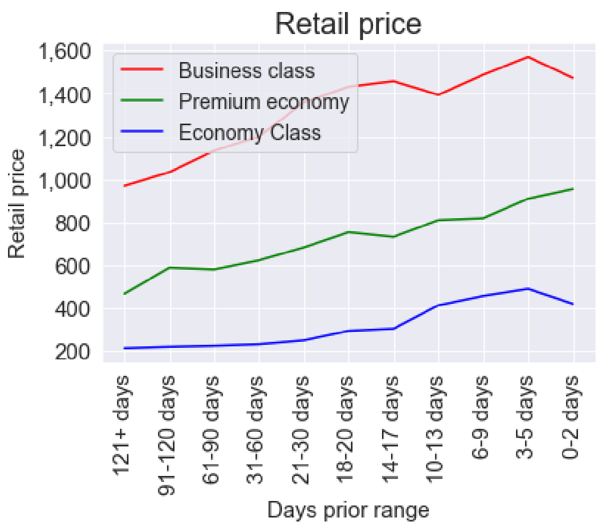

Notes: The horizontal axis shows the booking horizon by period before departure. The vertical axis displays the retail price paid to access economy, premium economy and business class. I exclude travelers belonging to frequent flyer programs and upgrade sales.

Figure 4: Prices

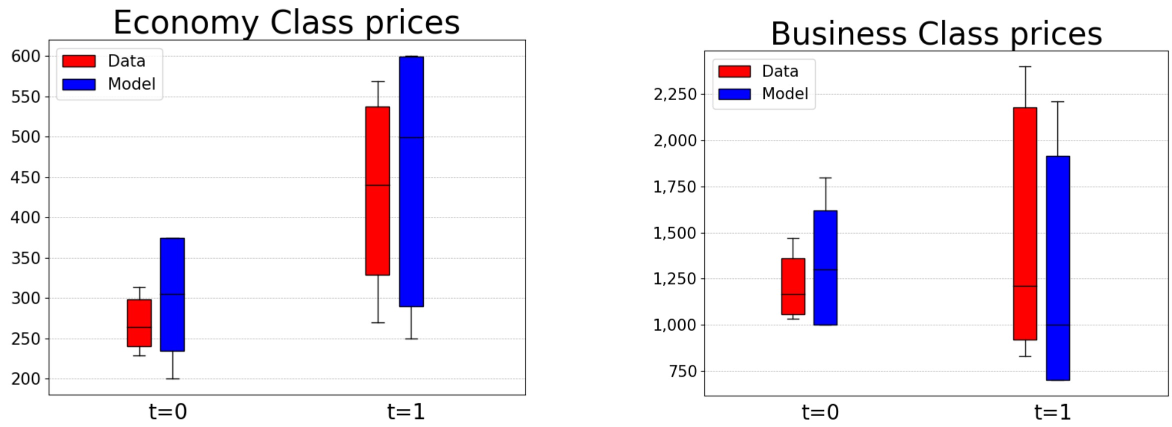

Notes: Price distribution for flights flying on the international route on aircraft with $c_{EC} = 247$ and $c_{BC} = 51$. The box represents the central 50% of the distribution with the median within it; the lower and upper bounds of the whisker represent the minimum and maximum of the distribution. The red box represents the actual distribution, whereas the blue one represents that for the simulated bootstrapped data.

Figure 5: Retail sales

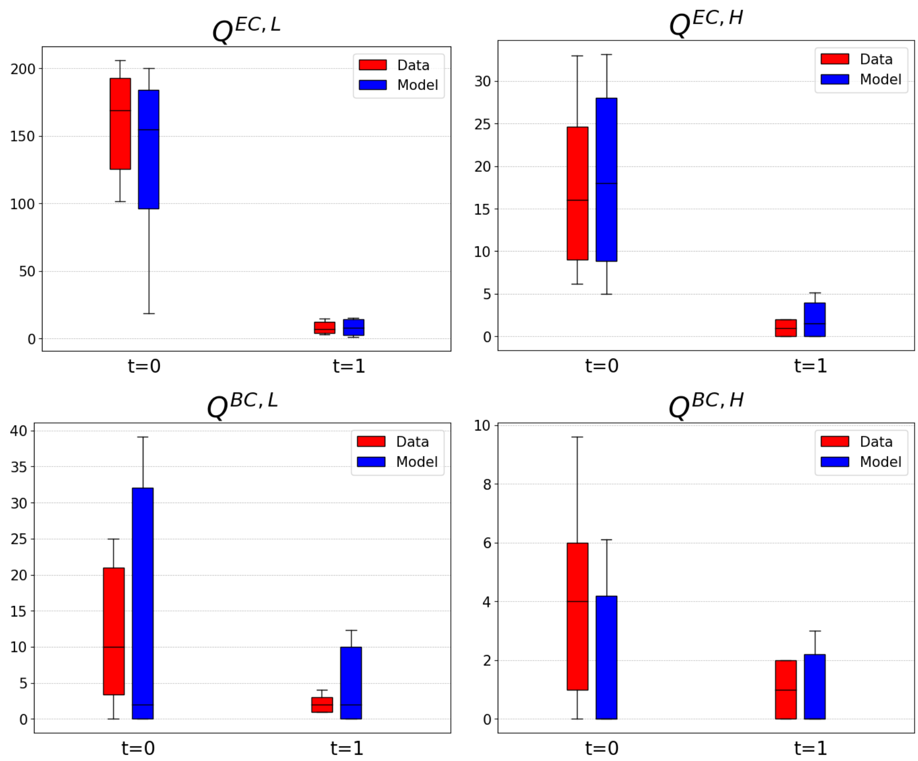

Notes: Distributions of simulated and actual retail purchases of tickets for flights on the international route on aircraft with $c_{EC} = 247$ and $c_{BC} = 51$. The box represents the central 50% of the distribution with the median within it; the lower and upper bounds of the whisker represent the minimum and maximum of the distribution. The red box represents the actual distribution, whereas the blue one represents that for the simulated bootstrapped data.

Figure 6: Upgrade prices and upgrade sales

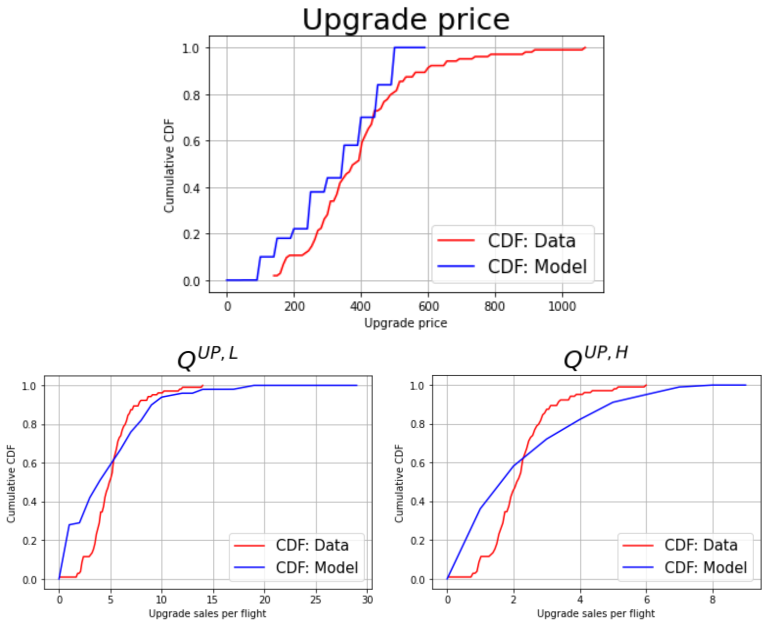

Notes: Distributions of simulated and actual prices and quantities of upgrades for flights on the international route on aircraft with $c_{EC} = 247$ and $c_{BC} = 51$. The red line represents the actual CDF of upgrade prices, whereas the blue line represents the average CDF of upgrade prices across the bootstrap estimates.

Figure 7: Counterfactual over time with capacity constraints

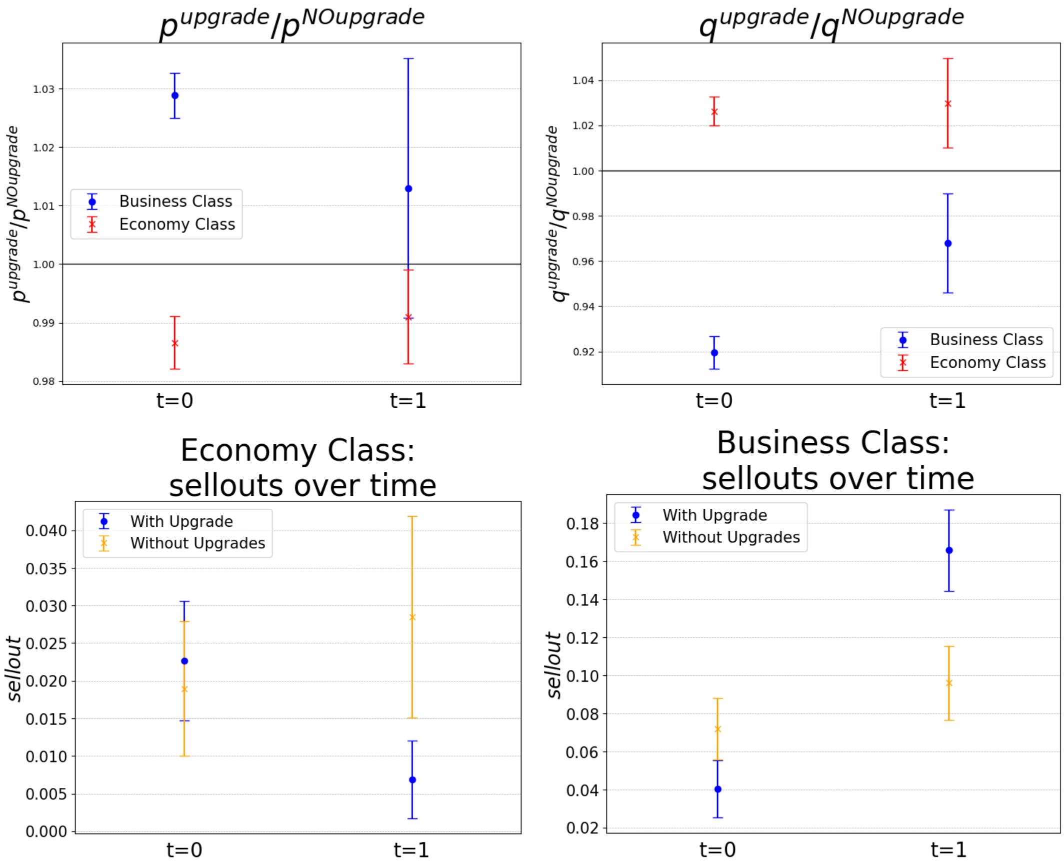

Notes: I evaluate the estimates and bootstrap 95% confidence intervals for large aircraft on the international route with capacity constraints. I use 10 bootstrap samples, each simulating 500 aircraft. In the sellouts panels, I report the fraction of flights experiencing sellouts in the corresponding cabin over time. Economy class sellouts decrease over time, as travelers upgrade upgrade to business class and leave economy class.

Figure 8: Counterfactual over time without capacity constraints

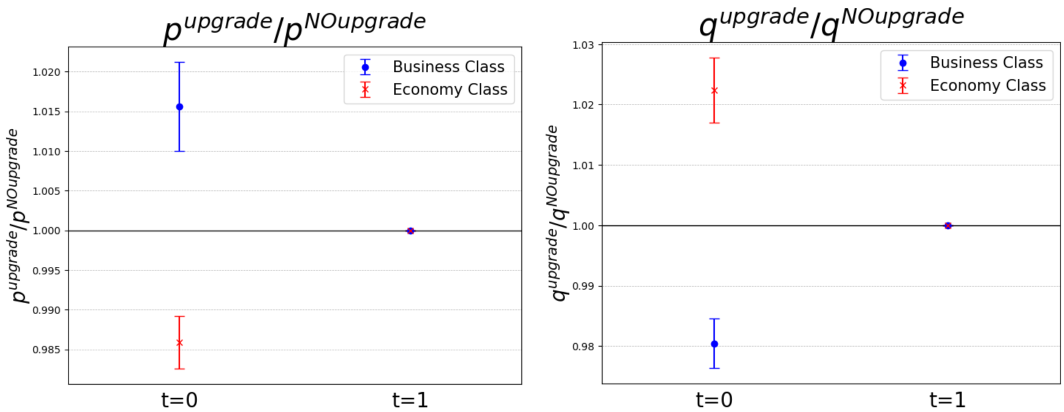

Notes: I evaluate the estimates and bootstrap 95% confidence interval for large aircraft on the international route without capacity constraints. I use 10 bootstrap samples, each simulating 500 aircraft.

A Appendix

A.1 Appendix – Example of an upgrade notification

Figure 9: Example of a notification email for an upgrade offer

Notes: Email notifying for an upgrade possibility

A.2 Appendix - Upgrade discount: Robustness check

The evidence from regression (1) remains robust even after I control for various confounding factors, including load factors and time and flight fixed effects. I summarize the results in Tables 18 and 19.

By defining t as the time of the initial purchase (so that $t = T - 2$ indicates the second-to-last day of the booking horizon when the airline allows upgrade sales) and f as the flight on which transaction i takes place, I estimate the following specification for robustness checks:

$$P_{\text{itf}} = \alpha + \beta_{\text{up}}\mathbb{I}\{\text{upgrade sale}\}_i + LF_{tf}\gamma_1 + LF_{T-2f}\gamma_2 + t\delta + FE_f + \epsilon_i \tag{12}$$

with LF_t indicating the load factor, defined as the ratio of realized sales up to period t to the total capacity. Similarly to in the analysis in Section 2.3, here, I estimate regression (12) separately for business and premium economy class. The results in Tables 18 and 19 show that, in all four specifications, the extent of the discount is statistically significant.

Low average final prices to access the premium cabin might occur on flights experiencing few premium-cabin sales. These flights might have low prices. Therefore the discounts observed in regression (1) might capture cross-flight price differences due to variation in unsold capacity rather than within-flight price differences. I account for this possibility in column (1), where I control for the load factor level at the time of the initial purchase. The price difference between retail and upgrade sales is still statistically different from 0.

Similarly, the airline might try to fill capacity in the premium cabin on those flights with an emptier premium cabin in the last two days before departure, when the upgrade program starts. Therefore, a lower level of LF_{T-2} might imply low premium-cabin retail prices, and thus, the observed discount from regression (1) might detect a load factor difference in the two days before departure across flights. This turns out not to be case, as the coefficient β_{UP} in column (2) is still statistically significant.

The difference between retail and upgrade final prices paid might be due to the initial ticket purchase date. In particular, as prices tend to increase over time, it might be that most upgrading customers have made their lower-cabin ticket purchase very early on in the booking horizon, when prices are typically low. This might imply low average final prices paid after an upgrade simply because of the low initial retail prices. I control for this effect in column (3) by including a time trend effect. The price difference between retail and upgrade sales remains statistically different from 0.

Finally, I jointly consider the load factor and time and flight fixed effects. I include flight fixed effects as upgrades might take place on specific flights that, for unobserved idiosyncratic reasons, might display low premium seat prices. Even with these controls, the price difference between retail and upgrade sales is still statistically different from 0.

Table 18: Business class

Variable	estimate (1)	estimate (2)	estimate (3)	estimate (4)
Average retail price (no upgrade)	1,040.30***	1068.92***	1,437.07***	1,150.04***
	(12.87)	(31.11)	(16.22)	(153.92)
Savings due to upgrade	-166.86***	-144.03***	-197.19***	-200.60***
	(14.35)	(14.18)	(15.47)	(19.86)
LF_t	yes	no	no	yes
LF_{T-2}	no	yes	no	no
t	no	no	yes	yes
Flight FE	no	no	no	yes

Notes: Results are in $ and computed from sales of business class seats. I drop sales to customers with flyers-IDs and sales with auction upgrades. Standard errors are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors are clustered at the flight level.

Table 19: Premium economy class

Variable	estimate (1)	estimate (2)	estimate (3)	estimate (4)
Average retail price (no upgrade)	526.49***	588.21***	786.88***	529.31***
	(8.30)	(16.20)	(5.80)	(12.13)
Savings due to upgrade	-112.24***	-115.76***	-149.61***	-111.13***
	(7.43)	(8.53)	(7.42)	(7.53)
LF_t	yes	no	no	yes
LF_{T-2}	no	yes	no	no
t	no	no	yes	yes
Flight FE	no	no	no	yes

Notes: Results are in $ and computed from sales of premium economy class seats. I drop sales to customers with flyers-IDs and sales with auction upgrades. Standard errors are in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors are clustered at the flight level.

A.3 Appendix – Estimation: The role of inattention parameters (α_L, α_H) with no capacity constraints

In this section, I analyze the effects of the upgrade option in the case with fully attentive economy class ticket holders when the airline does not face capacity constraints. The sequential price discrimination effect on the $t = 0$ customer is larger than that with inattentive travelers. In particular, since the probability that economy class ticket holders buy an upgrade increases, the opportunity cost of both raising business class prices and decreasing economy class prices decreases. Indeed, by charging higher business class and lower economy class prices, the number of economy class ticket buyers increase, thereby increasing the number of upgrading customers.

Table 20: Role of upgrades with attentive customers and without capacity constraints

	With Upgrades		Without Upgrades
	Attentive travelers	Inattentive travelers	
	$\alpha_H = \alpha_L = 1$	$\alpha_H = 0.2, \alpha_L = 0.1$	
	(1)	(2)	(3)
$p_{EC,0}$	282	313	317
	(2.8)	(2.95)	(2.49)
$p_{EC,1}$	430	430	430
	(5.28)	(5.28)	(5.28)
$p_{BC,0}$	1,238	1,152	1,134
	(12.56)	(14.23)	(11.96)
$p_{BC,1}$	848	848	848
	(15.61)	(15.61)	(15.61)
p_{UP}	242	268	
	(115)	(115)	
$passengers_{EC}$	98.09	145.69	150.27
	(1.83)	(2.25)	(2.82)
$passengers_{BC}$	106.68	42.36	35.13
	(1.77)	(1.73)	(1.61)
$upgrades$	78.62	7.78	
	(0.99)	(0.25)	

Notes: I evaluate the estimates and bootstrap standard errors for large aircraft on the international route when there are no capacity constraints under the case in which all travelers are attentive to the upgrade option (Column (1)) and under the case in which only the estimated fraction of travelers are attentive to the upgrade option (Column (2)). I use 10 bootstrap samples, each simulating 500 aircraft. The variable $passengers_k$ indicates the number of passengers flying in cabin k.

A.4 Appendix – Counterfactual benchmarks: Uniform pricing and free upgrades

In this section, to benchmark the welfare gains of the upgrade option within the literature, I consider two other scenarios: uniform pricing and free upgrades. The results in Table 21 show that, as the airline increases its pricing flexibility, both revenues and consumer surplus increase.

Table 21: Counterfactual – Various scenarios

	With Upgrades (1)	Without Upgrades (2)	Uniform Pricing (3)	Free Upgrades (4)
$p_{EC,0}$	303	308	317	308
	(3.33)	(2.95)	(2.58)	(2.95)
$p_{EC,1}$	418	422	317	422
	(5.29)	(5.64)	(2.58)	(5.64)
$p_{BC,0}$	1,305	1,268	1,156	1,268
	(12.77)	(11.32)	(6.1)	(11.32)
$p_{BC,1}$	996	983	1,156	983
	(13.92)	(24.15)	(6.1)	(24.15)
p_{UP}	297			0
	(134)			0
$passengers_{EC}$	164.5	167.2	168.74	153.55
	(2.85)	(3.24)	(2.86)	(3.23)
$passengers_{BC}$	29.3	23.82	20.51	37.47
	(0.86)	(0.83)	(0.84)	(0.66)
$upgrades$	7.11			13.64
	(0.19)			(0.26)
$sellout_{EC}$	0.01	0.03	0.04	0.01
	(0.004)	(0.009)	(0.011)	(0.003)
$sellout_{BC}$	0.166	0.096	0.095	0.354
	(0.016)	(0.015)	(0.014)	(0.022)
CS	81,978	80,751	80,156	83,645
	(3,017)	(3,049)	(3,016)	(3,013)
PS	80,451	78,897	77,823	78,897
	(2,251)	(2,214)	(2,188)	((2,214))
TS	162,430	159,649	157,980	162,544
	(5,244)	(5,239)	(5,177)	(5,203)

Notes: I evaluate the estimates and bootstrap standard errors for large aircraft on the international route. I use 10 bootstrap samples, each simulating 500 aircraft. The variable $passengers_k$ indicates the number of passengers flying in cabin k. The variable $sellouts_k$ indicates the fraction of flights (over the 500 simulated flights) that sell out in cabin k. **CS**, **PS** and **TS** indicate average (per flight) consumer, producer and total surplus, respectively.

The scenario of uniform pricing (Column (3)) considers a restrictive pricing regime where the airline sets a constant price over time for the two cabins in $t = 0$ after the realization of the initial demand shock and no upgrades are allowed. As the firm is constrained in its pricing decisions, producer surplus is lower than in the scenario with dynamic pricing but no upgrades (Column (2)). Relatively low business class prices in $t = 0$ lead to early sellouts in the premium cabin, whereas low economy class prices in

$t = 1$ lead to sellouts in the lower class cabin at the end of $t = 1$. Overall, early business class sellouts prevent high-value $t = 1$ customers from accessing high-quality products in $t = 1$, thereby leading to smaller consumer surplus with respect to that in the situation of dynamic pricing without upgrades. The results on prices, sellouts and total surplus align with the analysis of Williams (2022).

The second scenario involves the airline offering upgrades for free (Column (4)) while implementing dynamic pricing. With free upgrades, the upgrade option does not modify the opportunity cost of selling any seat in any period. Consequently, the airline pricing problem is the same as in the scenario without upgrades (Column (2)). As the airline cannot capture surplus from upgrading customers, upgrades benefit only customers. This leads to the highest levels of total surplus among all the counterfactual scenarios.

When the firm implements dynamic pricing, the introduction of the upgrade option increases producer surplus by 2%; transitioning from a uniform pricing regime to dynamic pricing results in a 1.3% increase in revenues. My model predicts that the upgrades are relatively more important in increasing revenues than the dynamic pricing. This result is primarily driven by the two-period assumption of the model. With more time periods, dynamic pricing would likely yield higher gains, as the airline could employ better intertemporal and intratemporal price discrimination strategies across customers with varying demands over time.[20] Additionally, with more than two periods, the scope of the upgrade option would be diminished, as the airline would aim at extracting surplus over the booking horizon through dynamic pricing rather than relying on last-day upgrades. Thus, the counterfactual results presented in this paper provide an upper bound on the effects of the upgrade option.

A.5 Appendix – Counterfactual results: The role of demand shocks

In this section, I complement the evidence in Tables 9, 10, and 11 regarding the effects of the upgrade option arising from various demand shocks by showing how different realizations of initial demand shocks lead to variation in upgrades sales, prices and surplus.

[20] As a benchmark in the literature, Williams (2022) finds that dynamic pricing increases revenues by 8% in a four-period model and one cabin.

Overall, as the demand shock increases, p_{UP} increases, and the number of upgrading travelers increases. Concurrently, producer and consumer surplus increase with the size of the demand shock.

Table 22: Leisure-type demand shock and upgrades

Demand Shock	Average p_{UP}
H,H	437.626
	(26.806)
H,M	353.364
	(13.033)
M,H	400.284
	(10.213)
H,L	201.505
	(32.788)
L,H	298.517
	(24.283)
M,M	286.710
	(6.956)
M,L	226.218
	(31.344)
L,M	257.299
	(8.939)
L,L	257.919
	(23.094)

Demand Shock	Average q_{UP}
H,H	6.132
	(1.134)
H,M	6.997
	(0.549)
M,H	6.662
	(0.383)
H,L	8.287
	(1.703)
L,H	6.275
	(0.850)
M,M	8.721
	(0.334)
M,L	6.228
	(0.650)
L,M	5.358
	(0.612)
L,L	3.370
	(0.518)

Notes: I simulate 500 flights operating on the international route under the scenario with upgrades. Demand shocks are the $t = 0$ leisure-type demand shocks, in the form of $\beta_{leisure,0}, \gamma_{leisure,0}$, where H represents a realization of the random coefficient in the top 20% of its distribution, L a realization in the bottom 20% and M a realization in the rest. I evaluate the optimal upgrade price and corresponding realized upgrade sales. Bootstrapped standard errors over 10 samples are recorded.

Table 23: Leisure-type demand shock and surplus

Demand Shock	Change in PS
H,H	1,844.5
	(1,092.2)
H,M	1,766.0
	(367.1)
M,H	1,521.7
	(475.4)
H,L	1,179.7
	(222.4)
L,H	2,362.9
	(650.4)
M,M	1,981.8
	(217.9)
M,L	835.0
	(122.6)
L,M	1,164.7
	(214.0)
L,L	534.7
	(235.8)

Demand Shock	Change in CS
H,H	1,641.1
	(1,737.9)
M,H	1,687.4
	(830.6)
H,M	430.9
	(717.7)
H,L	978.8
	(177.8)
L,H	4,098.5
	(1,011.3)
M,M	1,415.0
	(244.8)
M,L	838.0
	(306.5)
L,M	1,002.9
	(386.4)
L,L	331.9
	(269.7)

Notes: I simulate 500 flights operating on the international route under the scenarios with and without upgrades. Demand shocks are the $t = 0$ leisure-type demand shocks, in the form of $\beta_{leisure,0}, \gamma_{leisure,0}$, where H represents a realization of the random coefficient in the top 20% of its distribution, L a realization in the bottom 20% and M a realization in the rest. I evaluate the surplus in both scenarios and take the difference; for example, *Change in* $PS = PS^U - PS^{NoU}$, where PS^U is the producer surplus in the scenario with the upgrade option and PS^{NoU} is the producer surplus in the scenario without it. Bootstrapped standard errors over 10 samples are recorded.

A.6 Appendix – Estimation: The role of inattention

In this section, I show estimation results under the assumption that all customers pay attention to the email sent by the airline regarding the upgrade option. The focus is on the large aircraft flying on the international route.

The moments used for estimation are the same as those in Section 4.2, and the simulated model restricts $\alpha_H = \alpha_L = 1$ to assume fully attentive travelers. The estimation results are slightly different than those reported in Section 4.2, and visually, assuming fully attentive travelers leads to a worse model fit. When we assume fully attentive travelers, the model implies a larger number of upgrading customers, which modifies the airline's pricing strategies by inducing higher business class prices over the booking horizon with respect to those in the scenario with inattentive customers. Since sales are generated together with prices in the simulation routine, the demand estimates reported in Tables 24 and 25 are different from those in Tables 4 and 5.

Table 24: Preferences – International route, attentive customers

t=0		*t=1*	
parameter	*estimate*	*parameter*	*estimate*
(μ_{H0}, σ_{H0})	(314.24,94.24)	(μ_{H1}, σ_{H1})	(314.24,94.24)
	((45.64),(45.64))		((45.64),(45.64))
κ_{H0}	3.23	κ_{H1}	2.95
	(0.0)		(0.01)
(μ_{L0}, σ_{L0})	(324.39,193.08)	(μ_{L1}, σ_{L1})	(324.39,220.05)
	((8.72),(136.72))		((8.72),(25.0))
κ_{L0}	3.5	κ_{L1}	3.27
	(0.0)		(0.0)

Notes: Parameters for flights on the international route on aircraft with $c_{EC} = 247$, $c_{PE} = 21$, $c_{BC} = 30$. In the simulated airline problem, business class includes premium economy. I report bootstrap standard errors in parentheses. I use 100 bootstrap samples by resampling at the flight level.

Table 25: Arrival process and inattention - international route, attentive customers

parameter	*estimate*
λ_0	512.0
	(0)
λ_1	42.45
	(3.9)
θ_0	0.07
	(0.0)
θ_1	0.13
	(0.0)

Notes: Parameters for flights on the international route on aircraft with $c_{EC} = 247$, $c_{PE} = 21$, $c_{BC} = 30$. In the simulated airline problem, business class includes premium economy. I report bootstrap standard errors in parentheses. I use 100 bootstrap samples by resampling at the flight level.

Figure 10: Prices – International route, attentive customers

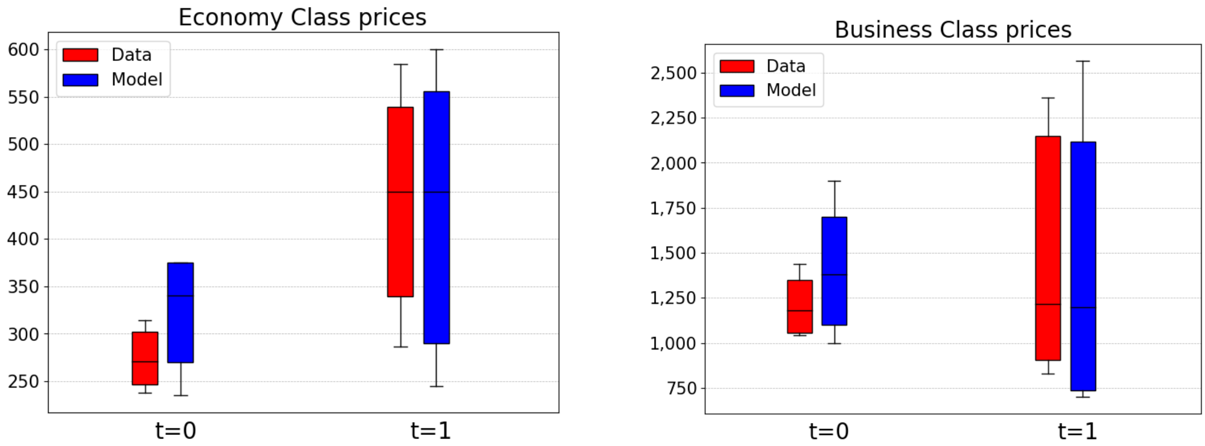

Notes: Simulated and actual price distributions for flights on the international route on aircraft with $c_{EC} = 247$, $c_{BC} = 51$, assuming fully attentive travelers ($\alpha_H = \alpha_L = 1$). The box represents the central 50% of the distribution with the median within it; the lower and upper bounds of the whisker represent the minimum and maximum of the distribution. The red box represents the actual distribution, whereas the blue one represents that for the simulated bootstrapped data.

Figure 11: Retail sales – International route, attentive customers

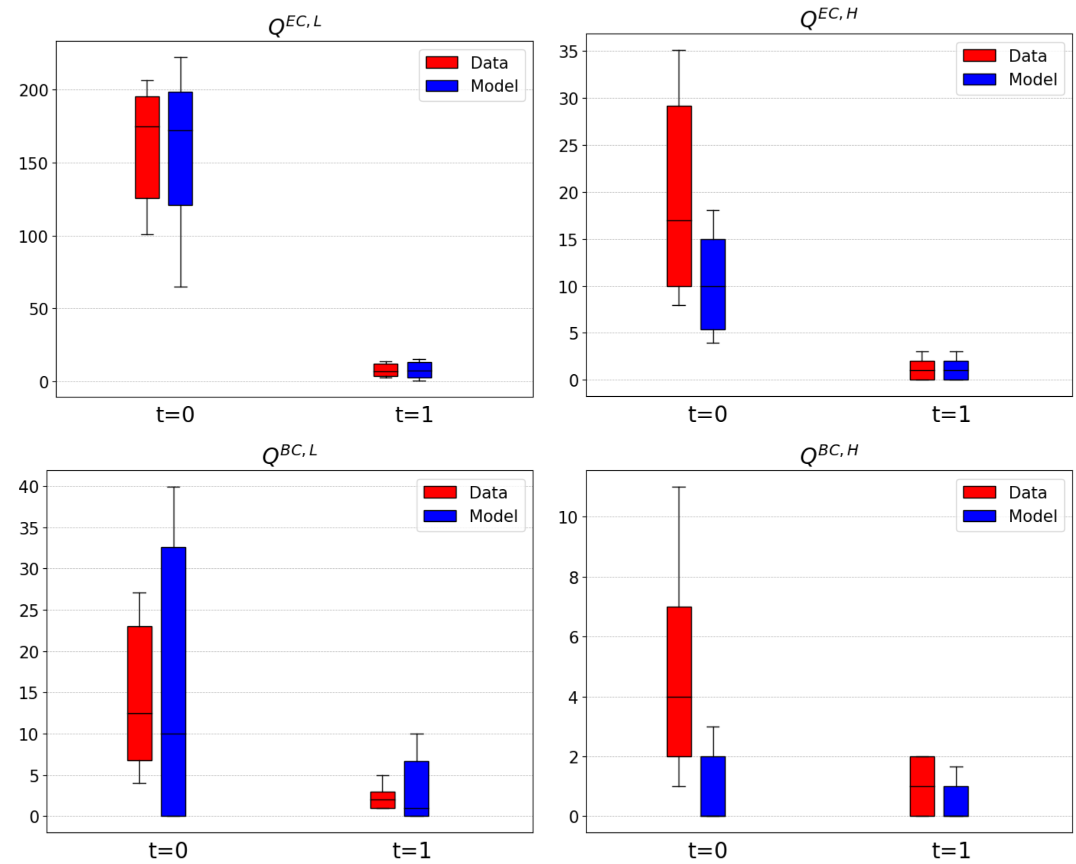

Notes: Simulated and actual price distributions for flights on the international route on aircraft with $c_{EC} = 247$, $c_{BC} = 51$, assuming fully attentive travelers ($\alpha_H = \alpha_L = 1$). The box represents the central 50% of the distribution with the median within it; the lower and upper bounds of the whisker represent the minimum and maximum of the distribution. The red box represents the actual distribution, whereas the blue one represents that for the simulated bootstrapped data.

Figure 12: Upgrade sales and prices – International route, attentive customers

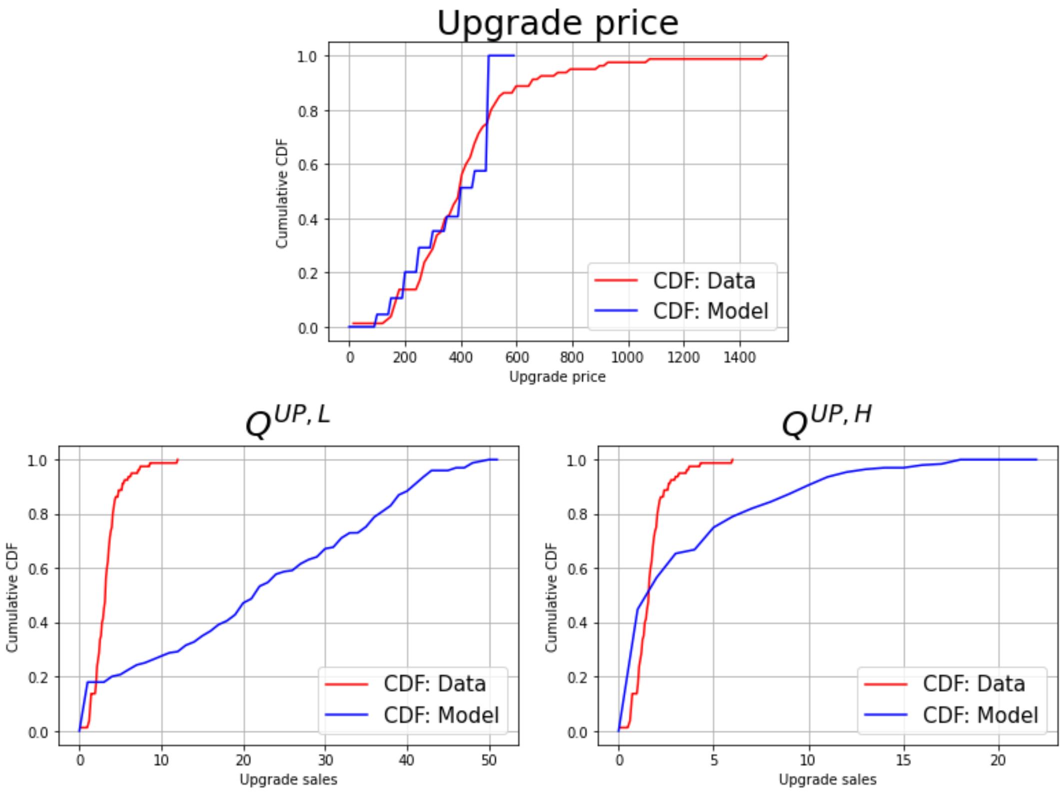

Notes: Distributions of simulated and actual price and quantities of upgrades for flights on the international route on aircraft with $c_{EC} = 247$, $c_{BC} = 51$, assuming fully attentive travelers ($\alpha_H = \alpha_L = 1$). The red line represents the actual CDF of upgrade prices, whereas the blue line represents the average CDF of upgrade prices across the bootstrap estimates.

A.7 Appendix – Small-size aircraft

A.7.1 Appendix – Distribution of aircraft size and upgrades

The number of upgrades varies based on the size of the aircraft. This paper does not delve into the discussion of how the airline deploys different aircraft on different routes; instead, I consider it to be exogenously determined.

The dataset used in this study consists of 11 types of aircraft, each varying in terms of capacity for economy, premium economy, and business class. The largest aircraft, which belongs to the Boeing 787 Dreamliner family, has a total of 247 seats in economy class (c_{EC}), with 21 seats in premium economy class (c_{PE}) and 30 seats in business class (c_{BC}). On the other hand, the smallest aircraft, the Boeing 757-200, does not have a premium economy class and features 153 seats in economy class (c_{EC}) and 16 seats in business class (c_{BC}). Analysis of the data, as shown in Table 26, reveals that the airline tends to deploy larger aircraft on the international route. Moreover, Table 1 illustrates that the majority of upgrades occur between economy and premium economy classes, leading to a higher number of upgrades on larger aircraft.

Table 26: Distribution of upgrades relative to aircraft size

	Domestic Route			International Route		
aircraft size	*upgrades*	*flights*	*upgrades per flight*	*upgrades*	*flights*	*upgrades per flight*
Large	58	15	3.8	1.932	294	6.5
Small	231	284	0.8	366	329	1.1

Notes: I report the distribution of the largest (labeled "Large", with $c_{EC} = 247, c_{PE} = 21, c_{BC} = 31$) and smallest aircraft (labeled "Small", with $c_{EC} = 153, c_{BC} = 16$) over the two routes. I also show the total number of upgrades by aircraft type.

A.7.2 Appendix – Estimation: Small aircraft

In this section, first, I present the estimation results on consumers' preferences, arrival rate, type mixture, and inattention coefficients for the smallest aircraft in the data, which operates on the domestic route. Then, I show how the estimated model fits the data.

The estimation results and fit for the small aircraft align with those of the larger aircraft.[21] Indeed, both aircraft have similar goodness of fit, according to the measure described in expression (11): the S^2 for small aircraft is 0.98 and that of the large one 0.975. However, due to the lower mean and variances of the observed prices on the domestic route, travelers exhibit lower mean and variances for both their willingness to pay and value for comfort than on the international route. Additionally, since the small aircraft has fewer seats, the expected arrival rate is lower than that for the international route. Notably, the estimates for the inattention coefficients are similar to those observed for the international route. By considering the average WTP and average prices, the simulated demand elasticities are -1.68 and -2.42 for business-type and leisure-type travelers, respectively.

Figures 13, 14 and 15 display how the model fits the observed prices, retail sales and

[21] In the original dataset, there are 284 small-sized aircraft flights on the domestic route. Similarly to in the case of large aircraft, discussed in Footnote 11, I estimate the model on a subset of flights to attain higher speed in the simulation algorithm and neglect outliers. I estimate the model on flights with $p_{EC,0} \in [200, 400]$, $p_{EC,1} \in [250, 650]$, $p_{BC,0} \in [300, 1100]$ and $p_{BC,1} \in [400, 1500]$. The final sample consists of 79 flights.

upgrades. With respect to the large aircraft on the international route, the simulated upgrade prices and sales for the small aircraft do not perfectly match the observed ones. This discrepancy arises from the fact that my model seldom predicts zero upgrade sales whereas zero upgrades are common in the actual data. Specifically, in the sample used for estimation of demand for small aircraft, the data show that almost 50% of flights do not have any upgrades.

Table 27: Preferences – Small aircraft on the domestic route

t=0		*t=1*	
parameter	*estimate*	*parameter*	*estimate*
(μ_{H0}, σ_{H0})	(319.6,119.5)	(μ_{H1}, σ_{H1})	(319.6,119.5)
	((2.81),(2.19))		((2.81),(2.19))
κ_{H0}	1.96	κ_{H1}	1.91
	(0.02)		(0.04)
(μ_{L0}, σ_{L0})	(230.0,140.0)	(μ_{L1}, σ_{L1})	(230.0,229.9)
	((0.0),(0.0))		((0.0),(1.0))
κ_{L0}	2.27	κ_{L1}	2.1
	(0.0)		(0.0)

Notes: Bootstrap estimates and standard errors for flights on the domestic route on aircraft with $c_{EC} = 153$ and $c_{BC} = 16$. The data used for estimation include 79 flights. I use 100 bootstrap samples, by resampling at the flight level.

Table 28: Arrival process and inattention – Small aircraft on the domestic route

parameter	*estimate*
λ_0	300
	(0.0)
λ_1	20.3
	(1.71)
θ_0	0.1
	(0.01)
θ_1	0.2
	(0.0)
α_H	0.19
	(0.01)
α_L	0.1
	(0.01)

Notes: Bootstrap estimates and standard errors for flights on the domestic route on aircraft with $c_{EC} = 153$ and $c_{BC} = 16$. The data used for estimation include 79 flights. I use 100 bootstrap samples, by resampling at the flight level.

Figure 13: Prices – Small aircraft on the domestic route

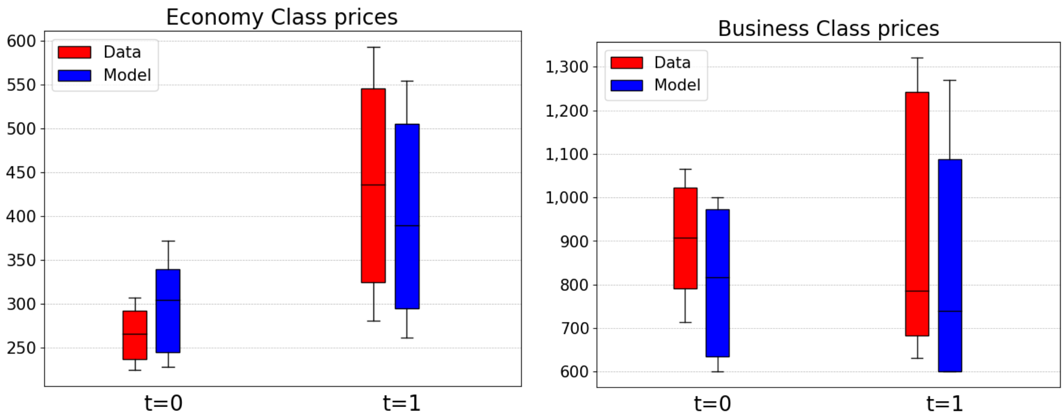

Notes: Simulated and actual price distributions for flights on the domestic route on aircraft with $c_{EC} = 153$ and $c_{BC} = 16$.

Figure 14: Retail sales – Small aircraft on the domestic route

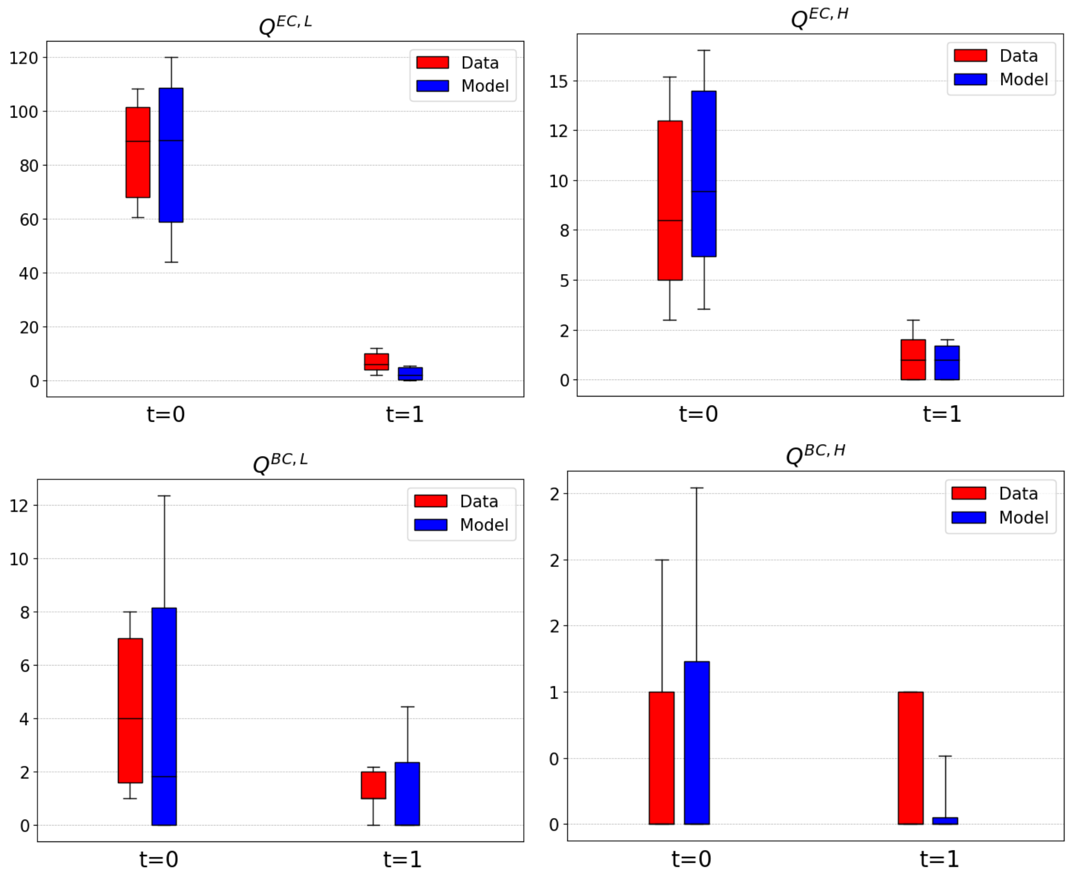

Notes: Simulated and actual price distributions for flights on the domestic route on aircraft with $c_{EC} = 153$ and $c_{BC} = 16$. The box represents the central 50% of the distribution with the median within it; the lower and upper bounds of the whisker represent the minimum and maximum of the distribution. The red box represents the actual distribution, whereas the blue one represents that for the simulated bootstrapped data.

Figure 15: Upgrade sales and prices – small aircraft on the domestic route

Notes: Distributions of simulated and actual price and quantities of upgrades for flights on the domestic route on aircraft with $c_{EC} = 153$ and $c_{BC} = 16$. The red line represents the actual CDF of upgrade prices, whereas the blue line represents the average CDF of upgrade prices across the bootstrap estimates.

A.7.3 Appendix - counterfactual: small aircraft

In this section, I assess the welfare consequences of the upgrade option in the context of small aircraft by counterfactual simulation. These results offer robustness for those presented in Section 5. First, I show that the upgrade option modifies the pricing strategy of the airline for small aircraft in a similar way as it does that for large aircraft. Then, I discuss the aggregate welfare consequences of upgrades and their distributional welfare consequences. Finally, I compare the impact of upgrades across aircraft of different sizes.

Economics of upgrades on small aircraft

Similarly to Table 6, Table 28 shows the effects of the introduction of the upgrade option with and without capacity constraints. In this way, I distinguish sequential price discrimination from inventory management. The scenario without capacity constraints illustrates how the upgrade option works as a way to sequentially price discriminate among $t = 0$ travelers. When the airline faces capacity constraints, upgrades work as a way to both price discriminate and manage inventory. In particular, with capacity constraints, the percentage change in $t = 0$ business class prices due to the introduction of the upgrade option is larger than that without capacity constraints. In this way, the airline reduces initial sellouts in business class. Higher $t = 0$ business class prices increase the number of travelers switching from the premium product to economy class. In turn, this reduces the airline's incentives to decrease initial economy class prices, so that the percentage decrease in economy class prices with capacity constraints is lower than that without them.

Table 28: Counterfactual in levels – Small aircraft on the domestic route

	WITH capacity constraints		WITHOUT capacity constraints	
	With upgrades (1)	Without upgrades (2)	With upgrades (3)	Without upgrades (4)
$p_{EC,0}$	287	289	318	322
	(4.35)	(3.9)	(4.69)	(4.67)
$p_{EC,1}$	401	406	423	423
	(5.8)	(6.25)	(6.92)	(6.92)
$p_{BC,0}$	955	919	798	783
	(9.46)	(10.53)	(8.23)	(8.99)
$p_{BC,1}$	805	785	715	715
	(10.08)	(7.71)	(6.08)	(6.08)
p_{UP}	204		141	
	(124)		(80)	
$passengers_{EC}$	88.35	90.02	66.64	68.1
	(1.72)	(1.67)	(1.58)	(1.62)
$passengers_{BC}$	9.45	7.09	21.25	18.67
	(0.22)	(0.27)	(1.05)	(1.04)
$upgrades$	3.34		3.5	
	(0.11)		(0.01)	

Notes: I simulate the estimates and bootstrap standard errors for small aircraft on the domestic route. I use 20 bootstrap samples, each simulating 500 aircraft. The scenario with capacity constraints considers $c_{EC} = 153$ and in business class is $c_{BC} = 16$. The variable $passengers_k$ indicates the number of passengers flying in cabin k.

Tables 29 and 30 show how the upgrade option works conditional on $t = 0$ leisure-type demand shocks. Similarly to Tables 7 and 8 for large aircraft, they show two facts. First, the airline's incentives to increase business class prices after the introduction of the upgrade option are stronger when the airline faces capacity constraints and a high demand shock (HH, HM, MH) is realized. In this way, the airline reduces initial-period retail sales and sellouts for the premium product. Second, when facing low demand shocks (LM, ML), the upgrade option works mainly as a way to sequentially price discriminate among $t = 0$ customers, and its effects with or without capacity constraints are similar since capacity constraints almost never bind.

Table 29: Price change, sequential price discrimination and inventory management – Small aircraft on the domestic route

Demand shock	WITH capacity constraints - Overall change $\%\Delta p_{0,EC}$	$\%\Delta p_{0,BC}$	WITHOUT capacity constraints - Change due to SPD $\%\Delta p_{0,EC}$	$\%\Delta p_{0,BC}$
H,H	-0.004	0.042	0.001	0.012
	(0.008)	(0.023)	(0.003)	(0.015)
H,M	0.010	0.065	0.000	0.028
	(0.005)	(0.017)	(0.000)	(0.010)
M,H	-0.011	0.103	-0.033	0.034
	(0.016)	(0.030)	(0.013)	(0.011)
H,L	-0.000	0.018	0.000	0.011
	(0.005)	(0.019)	(0.000)	(0.015)
L,H	-0.037	0.037	-0.040	0.039
	(0.024)	(0.024)	(0.022)	(0.022)
M,M	-0.004	0.062	-0.013	0.026
	(0.006)	(0.013)	(0.007)	(0.011)
M,L	0.000	0.011	-0.001	0.031
	(0.002)	(0.017)	(0.003)	(0.019)
L,M	-0.007	0.007	-0.008	0.008
	(0.006)	(0.006)	(0.008)	(0.008)
L,L	0.000	-0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)

Notes: I simulate 500 flights operating on the domestic route, and I report the percentage change induced by the introduction of the upgrade option. The overall change considers the percentage difference in $t = 0$ prices arising from the introduction of the upgrade option when the airline faces capacity constraints; thus, $\Delta p = \frac{p^u - p^{NOu}}{p^{NOu}}$, where p^u is the price when the upgrade option is available whereas p^{NOu} is the price when the upgrade option is not available. Columns under "change due to SPD" indicate the percentage change due to the introduction of the upgrade option when the airline does not face capacity constraints and, thus, associated with the use of upgrades as a sequential price discrimination (SPD) tool. Demand shocks are the $t = 0$ leisure-type demand shocks, in the form of $\beta_{leisure,0}, \gamma_{leisure,0}$, where H represents a realization of the random coefficient in the top 20% of its distribution, L a realization in the bottom 20% and M a realization in the rest. Bootstrapped standard errors over 20 samples are recorded.

Table 30: Quantity change, sequential price discrimination and inventory management – Small aircraft on the domestic route

Demand shock	WITH capacity constraints - Overall change		WITHOUT capacity constraints - Change due to SPD	
	$\%\Delta q_{0,EC}$	$\%\Delta q_{0,BC}$	$\%\Delta q_{0,EC}$	$\%\Delta q_{0,BC}$
H,H	0.023	-0.102	0.020	-0.023
	(0.012)	(0.056)	(0.032)	(0.029)
H,M	0.008	-0.106	0.039	-0.053
	(0.006)	(0.028)	(0.014)	(0.019)
M,H	0.074	-0.188	0.216	-0.076
	(0.027)	(0.041)	(0.087)	(0.021)
H,L	0.000	-0.009	0.001	-0.038
	(0.004)	(0.027)	(0.002)	(0.039)
L,H	0.215	-0.051	0.280	-0.057
	(0.147)	(0.042)	(0.176)	(0.038)
M,M	0.028	-0.090	0.071	-0.042
	(0.012)	(0.019)	(0.036)	(0.017)
M,L	0.002	-0.008	0.005	-0.009
	(0.005)	(0.015)	(0.006)	(0.011)
L,M	0.030	-0.000	0.042	-0.002
	(0.031)	(0.000)	(0.058)	(0.006)
L,L	0.000	0.000	0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)

Notes: I simulate 500 flights operating on the domestic route, and I report the percentage change induced by the introduction of the upgrade option. The overall change considers the percentage difference in $t = 0$ quantity arising from the introduction of the upgrade option when the airline faces capacity constraints; thus, $\Delta q = \frac{q^u - q^{NOu}}{q^{NOu}}$, where q^u is the quantity when the upgrade option is available whereas q^{NOu} is the quantity when the upgrade option is not available. Columns under "change due to SPD" indicate the percentage change due to the introduction of the upgrade option when the airline does not face capacity constraints and, thus, associated with the use of upgrades as a sequential price discrimination (SPD) tool. Demand shocks are the $t = 0$ leisure-type demand shocks, in the form of $\beta_{leisure,0}, \gamma_{leisure,0}$, where H represents a realization of the random coefficient in the top 20% of its distribution, L a realization in the bottom 20% and M a realization in the rest. Bootstrapped standard errors over 20 samples are recorded.

Tables 31, 32, 33, 34 and 35 show the effects of the introduction of the upgrade option for $t = 1$ prices, sales and sellouts on small aircraft flying on the domestic route. The results align with those for large aircraft flying on the international route in Tables 9, 10, 11, 22 and 23.

Table 31: Leisure-type demand shock and the effect of introducing upgrades on prices – Small aircraft on the domestic route

Demand shock	$\%\Delta p_{1,EC}$
H,H	0.011
	(0.040)
H,M	-0.003
	(0.020)
M,H	0.018
	(0.021)
H,L	-0.002
	(0.044)
L,H	0.026
	(0.051)
M,M	0.002
	(0.009)
M,L	-0.001
	(0.017)
L,M	0.009
	(0.020)
L,L	0.001
	(0.025)

Demand shock	$\%\Delta p_{1,BC}$
H,H	0.138
	(0.091)
H,M	0.076
	(0.038)
M,H	0.050
	(0.033)
H,L	0.021
	(0.048)
L,H	0.042
	(0.052)
M,M	0.066
	(0.016)
M,L	0.042
	(0.023)
L,M	0.046
	(0.035)
L,L	0.041
	(0.056)

Notes: I simulate 500 flights operating on the domestic route under two scenarios: with and without upgrades. In both cases, the airline faces the same demand shocks and capacity constraints. I evaluate optimal prices for any flight under the two scenarios and then consider their relative difference, in particular the *Change in* $p = \frac{p^U - p^{NoU}}{p^{NoU}}$, with p^U being the scenario with upgrades and p^{NoU} being the scenario without upgrades. Demand shocks are the $t = 0$ leisure-type demand shocks, in the form of $\beta_{leisure,0}, \gamma_{leisure,0}$, where H represents a realization of the random coefficient in the top 20% of its distribution, L a realization in the bottom 20% and M a realization in the rest. Bootstrapped standard errors over 20 samples are recorded.

Table 32: Leisure-type demand shock and the effect of introducing upgrades on retail sales – Small aircraft on the domestic route

Demand Shock	$\%\Delta q_{1,EC}$
H,H	0.149
	(0.173)
H,M	0.153
	(0.080)
M,H	0.099
	(0.086)
H,L	0.087
	(0.172)
L,H	-0.004
	(0.123)
M,M	0.103
	(0.052)
M,L	0.100
	(0.085)
L,M	0.037
	(0.091)
L,L	0.040
	(0.086)

Demand Shock	$\%\Delta q_{1,BC}$
H,H	-0.123
	(0.099)
H,M	-0.063
	(0.073)
M,H	-0.029
	(0.075)
H,L	0.011
	(0.078)
L,H	0.066
	(0.096)
M,M	-0.055
	(0.038)
M,L	-0.002
	(0.058)
L,M	-0.010
	(0.053)
L,L	-0.005
	(0.108)

Notes: I simulate 500 flights operating on the domestic route under two scenarios: with and without upgrades. In both cases, the airline faces the same demand shocks. Demand shocks are the $t = 0$ leisure-type demand shocks, in the form of $\beta_{leisure,0}, \gamma_{leisure,0}$, where H represents a realization of the random coefficient in the top 20% of its distribution, L a realization in the bottom 20% and M a realization in the rest. I evaluate retail sales for any flight under the two scenarios and then consider their percentage difference, in particular *Change in q* $= \frac{q^U - q^{NoU}}{q^{NoU}}$, with q^U being the quantity in the scenario with upgrades, and q^{NoU} the quantity in the scenario without upgrades. Bootstrapped standard errors over 20 samples are recorded.

Table 33: Leisure-type demand shock and the effect of introducing upgrades on sellouts – Small aircraft on the domestic route

Demand Shock	$\%\Delta sellout_{0,EC}$	Demand Shock	$\%\Delta sellout_{0,BC}$
H,H	0.00781	H,H	-0.10278
	(0.02034)		(0.05829)
H,M	0.00000	H,M	-0.08073
	(0.00000)		(0.04807)
M,H	0.00000	M,H	-0.13829
	(0.00000)		(0.03135)
H,L	0.00000	H,L	0.00000
	(0.00000)		(0.00000)
L,H	0.00000	L,H	0.00000
	(0.00000)		(0.00000)
M,M	0.00024	M,M	-0.03197
	(0.00108)		(0.00969)
M,L	0.00000	M,L	0.00000
	(0.00000)		(0.00000)
L,M	0.00000	L,M	0.00000
	(0.00000)		(0.00000)
L,L	0.00000	L,L	0.00000
	(0.00000)		(0.00000)

Demand Shock	$\%\Delta sellout_{1,EC}$	Demand Shock	$\%\Delta sellout_{1,BC}$
H,H	-0.00574	H,H	0.18410
	(0.03102)		(0.08090)
H,M	-0.00829	H,M	0.17803
	(0.01130)		(0.06451)
M,H	0.00000	M,H	0.04519
	(0.00000)		(0.05242)
H,L	-0.04520	H,L	0.05660
	(0.05046)		(0.03717)
L,H	0.00000	L,H	0.00870
	(0.00000)		(0.02517)
M,M	-0.00026	M,M	0.11107
	(0.00118)		(0.02889)
M,L	0.00000	M,L	0.01086
	(0.00000)		(0.01523)
L,M	0.00000	L,M	0.01111
	(0.00000)		(0.01115)
L,L	0.00000	L,L	0.00955
	(0.00000)		(0.01999)

Notes: I simulate 500 flights operating on the domestic route under the two scenarios: with and without upgrades. In both cases, the airline faces the same demand shocks. Demand shocks are the $t = 0$ leisure-type demand shocks, in the form of $\beta_{leisure,0}, \gamma_{leisure,0}$, where H represents a realization of the random coefficient in top 20% of its distribution, L a realization in the bottom 20% and M a realization in the rest. I evaluate the fraction (over the 500 simulations) of flights that sell out; in particular *Change in sellout*$_k$ = $sellout_k^U - sellout_k^{NoU}$, with $sellout_k^U$ being the fraction of flights that sell out in cabin k in the scenario with the upgrade option, and $sellout_k^{NoU}$ being the fraction of flights that sellout in cabin k in the scenario without the upgrade option. Bootstrapped standard errors over 20 samples are recorded.

Table 34: Leisure-type demand shock and upgrades – Small aircraft on the domestic route

Demand Shock	Average p_{UP}
H,H	385.8
	(26.417)
H,M	260.0
	(14.105)
M,H	275.6
	(13.978)
H,L	141.9
	(15.661)
L,H	173.6
	(22.838)
M,M	181.9
	(6.478)
M,L	172.7
	(15.421)
L,M	167.2
	(18.929)
L,L	191.7
	(30.638)

Demand Shock	Average q_{UP}
H,H	3.1
	(0.625)
H,M	3.5
	(0.252)
M,H	3.6
	(0.230)
H,L	3.4
	(0.475)
L,H	2.7
	(0.592)
M,M	4.2
	(0.234)
M,L	2.7
	(0.224)
L,M	2.1
	(0.297)
L,L	1.3
	(0.279)

Notes: I simulate 500 flights operating on the domestic route under the scenario with upgrades. Demand shocks are the $t = 0$ leisure-type demand shocks, in the form of $\beta_{leisure,0}, \gamma_{leisure,0}$, where H represents a realization of the random coefficient in the top 20% of its distribution, L a realization in the bottom 20% and M a realization in the rest. I evaluate the optimal upgrade price and corresponding realized upgrade sales. Bootstrapped standard errors over 20 samples are recorded.

Table 35: Leisure-type demand shock and surplus – Small aircraft on the domestic route

Demand Shock	Change in PS
H,H	709
	(562.154)
H,M	558
	(189.943)
M,H	117
	(238.668)
H,L	267
	(109.010)
L,H	510
	(125.451)
M,M	381
	(60.948)
M,L	284
	(57.854)
L,M	253
	(44.803)
L,L	173
	(79.297)

Demand Shock	Change in CS
H,H	1,000
	(737.058)
H,M	234
	(206.631)
M,H	537
	(496.190)
H,L	241
	(179.311)
L,H	491
	(197.730)
M,M	440
	(176.876)
M,L	194
	(81.951)
L,M	184
	(49.300)
L,L	108
	(69.081)

Notes: I simulate 500 flights operating in the domestic route under the scenarios with and without upgrades. Demand shocks are the $t = 0$ leisure-type demand shocks, in the form of $\beta_{leisure,0}, \gamma_{leisure,0}$, where H represents a realization of the random coefficient in the top 20% of its distribution, L a realization in the bottom 20% and M a realization in the rest. I evaluate surplus in both scenarios and take the difference, for example *Change in* $PS = PS^U - PS^{NoU}$, with PS^U is the producer surplus in the scenario with the upgrade option and PS^{NoU} is the producer surplus in the scenario without it. Bootstrapped standard errors over 20 samples are recorded.

Aggregate welfare effects of the upgrade option

Table 36 compares the average effect of the upgrade option in terms of welfare. Then, I analyze the effects of the upgrade option over time and present them in Figure 16. The results are similar to those in Table 12 and Figure 7.

Table 36: Counterfactual in levels – Small aircraft on the domestic route

	With Upgrades	Without Upgrades	$\Delta = With - Without$ upgrades
$p_{EC,0}$	287	289	-2
	(4.35)	(3.9)	(1.08)
$p_{EC,1}$	401	406	-5
	(5.8)	(6.25)	(2.58)
$p_{BC,0}$	955	919	36
	(9.46)	(10.53)	(3.58)
$p_{BC,1}$	805	785	20
	(10.08)	(7.71)	(9.42)
p_{UP}	204		
	(124)		
$passengers_{EC}$	88.35	90.02	-1.672
	(1.72)	(1.67)	(0.35)
$passengers_{BC}$	9.45	7.09	2.365
	(0.22)	(0.27)	(0.15)
upgrades	3.34		
	(0.11)		
$sellout_{EC}$	0.0	0.01	-0.0
	(0.002)	(0.004)	(0.003)
$sellout_{BC}$	0.201	0.122	0.079
	(0.019)	(0.014)	(0.017)
CS	31,274	30,919	355
	(1,157)	(1,157)	(100)
PS	32,319	31,968	351
	(940)	(942)	(60)
TS	63,594	62,887	706
	(2,076)	(2,077)	(124)

Notes: I evaluate the estimates and bootstrap standard errors for small aircraft on the domestic route. I use 20 bootstrap samples, each simulating 500 aircraft. The variable $passengers_k$ indicates the number of passengers flying in cabin k. The variable $sellouts_k$ indicates the fraction of flights (over the 500 simulated flights) that sell out in cabin k. **CS**, **PS** and **TS** indicate average (per flight) consumer, producer and total surplus respectively.

Figure 16: Counterfactual over time – Small aircraft on the domestic route

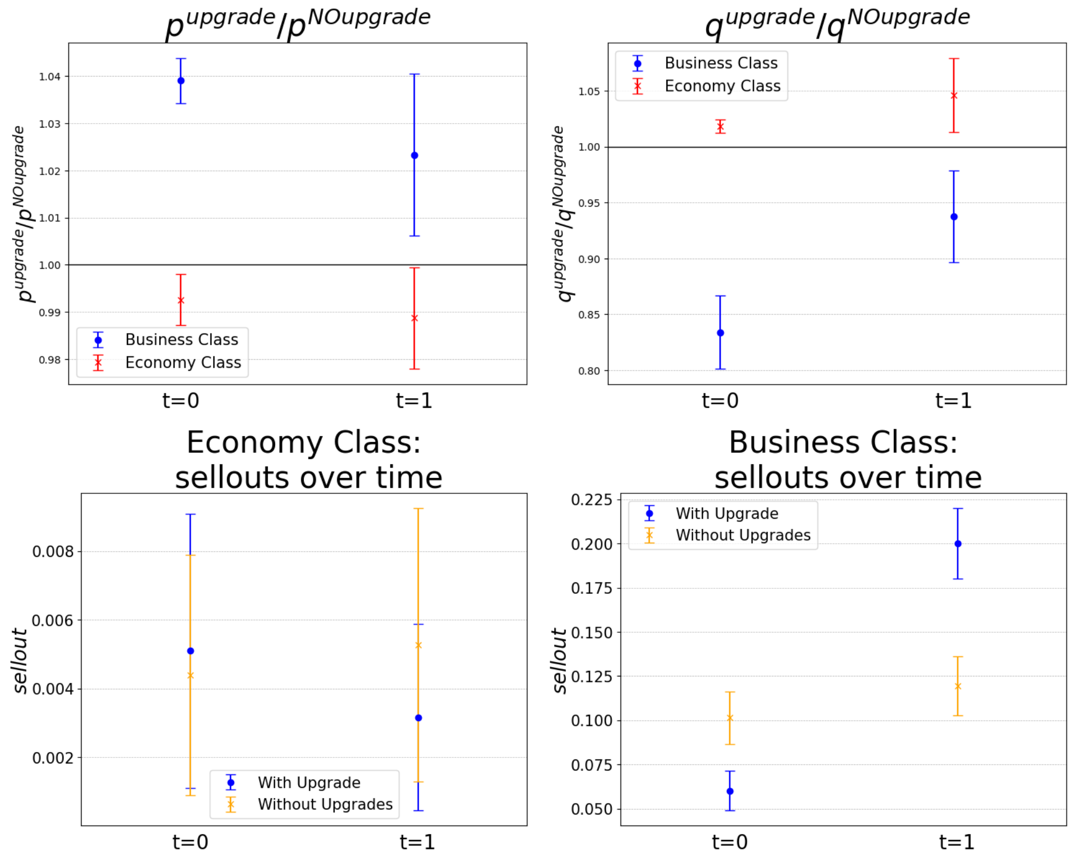

Notes: I evaluate estimates and bootstrap 95% confidence intervals for means for small aircraft on the domestic route. I use 20 bootstrap samples, each simulating 500 aircraft.

Distributional welfare effects of the upgrade option

Tables 37, 38 and 39 show how customers' behavior and welfare change according to the new pricing decisions arising from the introduction of the upgrade option. The results align with those for the international route reported in Section 5.2.

Table 37: Counterfactual, change in the absolute number of retail sales – Small aircraft on the domestic route

	Eliminating upgrades...	**Without Upgrades** Outside Option	*EC*	*BC*	*total*
With Upgrades	Outside Option	114,814	478	15	115,307
		(1,044)	(113)	(4)	(1043)
	EC	804	42,756	570	44,130
		(138)	(818)	(75)	(862)
	$EC + UP$	14	1,598	50	1,662
		(3)	(50)	(8)	(53)
	BC	41	110	2,897	3,048
		(12)	(20)	(107)	(113)
	total	115,673	44,942	3,532	
		(995)	(837)	(138)	

Notes: I simulate 500 flights under the two scenarios: with and without upgrades for 20 bootstrap samples. I report bootstrap estimates and standard errors. I evaluate how the behavior of consumers changes when the upgrade option is eliminated.

Table 38: Counterfactual, change in retail sales as percentage of original purchases – Small aircraft on the domestic route

	Eliminating upgrades...	**Without Upgrades** Outside Option	*EC*	*BC*
With Upgrades	Outside Option	99.6%	0.4%	0.0%
		(0.1)	(0.1)	(0.0)
	EC	1.8%	96.8%	1.3%
		(0.3)	(0.4)	(0.2)
	$EC + UP$	0.8%	96.1%	3.0%
		(0.2)	(0.5)	(0.5)
	BC	1.4%	3.6%	95.0%
		(0.4)	(0.6)	(0.8)

Notes: I simulate 500 flights under the two scenarios with and without upgrades for 20 bootstrap samples. I report bootstrap estimates and standard errors. I evaluate how the behavior of consumers changes in percentage terms when the upgrade option is eliminated.

Table 39: Counterfactual, consumer surplus net effects – Small aircraft on the domestic route

Eliminating Upgrades ↗	Outside Option	*EC*	*BC*
Outside Option	0	-17,073	-1,366
	(0)	(4,929)	(876)
EC	57,163	36,743	-40,394
	(14,443)	(38,522)	(10,247)
$EC + UP$	1,443	132,920	9,537
	(663)	(4,550)	(2,176)
BC	47,135	1,589	-50,081
	(17,723)	(2,389)	(10,026)

Notes: I simulate 500 flights under the scenarios with and without upgrades for 20 bootstrap samples. I report bootstrap estimates and standard errors. I evaluate consumer surplus for any passenger on all 500 flights under the two scenarios and then consider their difference: $CS^{\text{scenario with upgrades}} - CS^{\text{scenario without upgrades}}$. Results are in \$.

Tables 40 and 41 show how the introduction of upgrades reduces business class retail sales while at the same time increasing total revenues, thanks to an increase in revenues from economy class and upgrade fees. The results are similar to those discussed in Section 5.2.

Table 40: Producer surplus counterfactual, aggregate revenues across cabins – Small aircraft on the domestic route

Cabin	With Upgrades	Without Upgrades	Δ
EC	26.768\$	26.357\$	+411\$
	(0.734)	(0.725)	(62)
BC	4.977\$	5.612\$	-635\$
	(0.248)	(0.268)	(103)
UP	575\$	0\$	+575\$
	(20)	0	
total	32.320\$	31.969\$	+351\$
	(0.917)	(0.919)	(59.0)

Notes: I simulate 500 flights under the scenarios: with and without upgrades for 20 bootstrap samples. I report bootstrap estimates and standard errors. I evaluate producer surplus for any flight under the scenarios with and without upgrades; then, I consider the average across the 500 flights. Results are in \$.

Table 41: Producer surplus counterfactual, distribution of revenues across cabins – Small aircraft on the domestic route

Cabin	With Upgrades	Without Upgrades
EC	0.828	0.825
	(0.005)	(0.005)
BC	0.154	0.175
	(0.005)	(0.005)
UP	0.018	0
	(0.001)	(0)

Notes: I simulate 500 flights under the scenarios with and without upgrades for 20 bootstrap samples. I report bootstrap estimates and standard errors. I evaluate the average distribution of producer surplus across products by scenario.

Comparison across aircraft sizes

Table 42 compares the percentage changes in relevant outcomes resulting from the introduction of the upgrade option across aircraft with different levels of capacity. The upgrade option has similar effects across aircraft of different sizes. However, in terms of producer surplus, the upgrade option generates relatively more incremental revenues on large aircraft. This is due to the fact that, when there is a small demand shock along the domestic route, the airline does not use the upgrade option to sequentially price discriminate among $t = 0$ travelers, as shown in Table 29. With respect to the international route, two factors likely explain this result. On one hand, with respect to travelers on the international route, domestic travelers have a lower value for comfort. On the other hand, the prices of economy class and upgrades are similar across the domestic and international routes. Therefore, when there is a low realization of demand, the share of economy class ticket holders at the margin for business class have smaller taste for comfort and are then less likely to purchase an upgrade. Therefore, the upgrade option does not affect the pricing problem of the airline in case of small demand shocks.

Table 42: Counterfactual – Comparison across different aircraft

	$\%\Delta = \frac{X^{\text{With Upgrades}} - X^{\text{Without Upgrades}}}{X^{\text{Without Upgrades}}}$	
Outcome	Small Aircraft	Large Aircraft
$p_{EC,0}$	-0.007	-0.016
	(0.0)	(0.0)
$p_{EC,1}$	-0.012	-0.009
	(0.01)	(0.01)
$p_{BC,0}$	0.039	0.029
	(0.0)	(0.0)
$p_{BC,1}$	0.024	0.013
	(0.01)	(0.01)
$passengers_{EC}$	-0.019	-0.016
	(0.0039)	(0.0028)
$passengers_{BC}$	0.334	0.23
	(0.0332)	(0.0114)
*fraction of upgrading passengers**	0.05	0.05
	(0.0045)	(0.0023)
$sellout_{EC}$	-0.478	-0.7482
	(0.4438)	(0.136)
$sellout_{BC}$	0.6582	0.7453
	(0.1818)	(0.1593)
CS	0.012	0.015
	(0.002)	(0.002)
PS	0.011	0.020
	(0.002)	(0.002)
TS	0.011	0.017
	(0.001)	(0.002)

Notes: I evaluate the percentage change from introducing the upgrade option for small and large aircraft over 500 flights. For example, the variable $sellouts_k$ considers the percentage change due to the introduction of the upgrade option in the fraction of sellouts in cabin k. Row * considers the fraction of customers in economy class in $t = 0$ who upgrade to business class. Bootstrapped standard errors over 10 samples are recorded.

A.8 Appendix – Counterfactual: The role of capacity constraints

In this section, to understand how the upgrade option functions, I consider the interaction between the upgrade option and capacity constraints, following the discussion in Section 5.1. Without capacity constraints, the upgrade option works as a way to sequentially price discriminate among $t = 0$ customers. With capacity constraints, upgrades also serve as a way to manage inventory. To support my claims, first, I show the effect of the introduction of the upgrade option on average across all simulated flights, and then, I show how the results change based on demand shocks.

A.8.1 Appendix – Counterfactual: No capacity constraints

In this section, I focus on the role of upgrades in the absence of capacity constraints on average across all flights. When there are no capacity constraints, upgrades serve as a tool to sequentially price discriminate among $t = 0$ customers, which affects $t = 0$ prices. This increases welfare for both travelers and the firm.

In the absence of capacity constraints, introducing upgrades leads to an increase in business class prices and a decrease in economy class prices in $t = 0$, with no impact on $t = 1$ prices. The pricing decisions in one period do not affect the decisions in the other period. This happens because the intertemporal connection between prices lies in the evolution of remaining capacity over time but, in this scenario, there are no capacity constraints. In this scenario, upgrades serve as a tool for sequential price discrimination, allowing the airline to implement second-degree price discrimination twice. In period $t = 0$, it screens all customers between higher- and lower-quality products; then, in period $t = 1$, it sorts lower-quality product holders into those who choose to upgrade and those who do not. The effects of the upgrade option on $t = 0$ prices are similar to the effects of introducing an intermediate-quality good between economy and business class in period $t = 0$. This concept is similar to the idea of "versioning" discussed by Varian (1989), where a new version of a product is marketed at an intermediate price, the overall number of customers increases and the welfare of both consumers and the seller rises. Similarly, as shown in Table 43, the upgrade option increases the number of travelers flying and welfare for both customers and travelers.

Table 43: Counterfactual, role of capacity constraints – Large aircraft on the international route

	WITH capacity constraints		WITHOUT capacity constraints	
	With Upgrades (1)	Without Upgrades (2)	With Upgrades (3)	Without Upgrades (4)
$p_{EC,0}$	303	308	313	317
	(3.33)	(2.95)	(2.95)	(2.49)
$p_{EC,1}$	418	422	430	430
	(5.29)	(5.64)	(5.28)	(5.28)
$p_{BC,0}$	1,305	1,268	1,152	1,134
	(12.77)	(11.32)	(14.23)	(11.96)
$p_{BC,1}$	996	983	848	848
	(13.92)	(24.15)	(15.61)	(15.61)
p_{UP}	297		268	
	(134)		(115)	
$passengers_{EC}$	164.5	167.2	145.69	150.27
	(2.85)	(3.24)	(2.25)	(2.82)
$passengers_{BC}$	29.3	23.82	42.36	35.13
	(0.86)	(0.83)	(1.73)	(1.61)
upgrades	7.11		7.78	
	(0.19)		(0.25)	
$sellout_{EC}$	0.01	0.03	0	0
	(0.004)	(0.009)	(0.0)	(0.0)
$sellout_{BC}$	0.166	0.096	0	0
	(0.016)	(0.015)	(0.0)	(0.0)
CS	81,978	80,751	87,248	85,897
	(3,017)	(3,049)	(3,655)	(3,675)
PS	80,451	78,897	84,414	82,536
	(2,251)	(2,214)	(2,587)	(2,558)
TS	162,430	159,649	171,663	168,434
	(5,244)	(5,239)	(6,216)	(6,211)

Notes: I evaluate the estimates and bootstrap standard errors for large aircraft on the international route. I use 10 bootstrap samples, each simulating 500 aircraft. The variable $passengers_k$ indicates the number of passengers flying in cabin k. The variable $sellouts_k$ indicates the fraction of flights (over the 500 simulated flights) that sell out in cabin k. **CS**, **PS** and **TS** indicate average (per flight) consumer, producer and total surplus, respectively.

A.8.2 Demand shock without capacity constraints

In this section, I complement the evidence from Table 7 on the consequences of the upgrade option without capacity constraints, conditional on the demand shock. I consider the effects of the introduction of the upgrade option. The results are similar to those in Table

10: as the demand shocks increase, upgrade prices increase, together with the number of upgrade sales.

Table 44: Leisure-type demand shock and upgrades – Large aircraft on the international route

Demand shock	Average p_{UP}
H,H	355.404
	(22.339)
H,M	282.196
	(11.118)
M,H	334.264
	(10.015)
H,L	191.543
	(26.889)
L,H	299.346
	(30.034)
M,M	267.218
	(7.772)
M,L	209.548
	(22.396)
L,M	259.447
	(13.379)
L,L	253.358
	(38.295)

Demand shock	Average q_{UP}
H,H	8.994
	(1.190)
H,M	11.199
	(0.401)
M,H	6.944
	(0.269)
H,L	9.272
	(1.699)
L,H	6.385
	(0.852)
M,M	8.594
	(0.346)
M,L	6.304
	(0.623)
L,M	5.484
	(0.609)
L,L	3.619
	(0.537)

Notes: I simulate 500 flights operating on the international route under the scenario without capacity constraints. I evaluate optimal upgrade prices for any flight and then simulate the corresponding demand for upgrades. Results are in \$, and demand shocks are in the form of $\beta_{leisure,0}, \gamma_{leisure,0}$. Bootstrapped standard errors are recorded.