A Market for Airport Slots

Marleen Marra*

March 6, 2024

Abstract

This paper assesses the welfare implications of introducing airport slot auctions. Given the absence of an existing market and available bids, slot values are derived from airlines' additional profits, utilizing a flight-level structural model. This model considers various factors, such as consumer preferences, scheduling efficiencies, and aircraft-specific costs, with a new flight-level dataset facilitating its estimation. The estimated network benefits carry through the market-based allocation, skewing the resulting slot allocation towards dominant carriers within their hub airports. Consumers benefit from the auction and surprisingly prefer a mechanism without quantity caps. This preference is attributed to the convenience benefits of larger carriers outweighing the gains from increased competition.

JEL: L13, L93, D44, C57

Keywords: *Airline industry, Structural estimation, Airport slots, Auctions*

**Sciences Po, 28 Rue des Saints-Pères, 75007 Paris France, e-mail: marleen.marra@sciencespo.fr.* I would like to thank Kevin Guitet, Sophie Berthon, and Thierry Delort from the French Civil Aviation Authority (DGAC) for providing me with the administrative data and for answering my numerous questions. I am also grateful to Nicolas Astier, Christian Bontemps, Moshe Buchinsky, Clément de Chaisemartin, Peter Cramton, Allan Crawford, Pierre Dubois, Junnan He, Michele Fioretti, Jan de Loecker, Myrto Kalouptsidi, Thierry Mayer, Mateusz Mysliwski, Aureo de Paula, Mar Reguant, Kevin Renmy, David Salant, Frank Verboven, and Kevin Williams for helpful discussions, as well as to other colleagues and to seminar participants at the University of California, Berkeley, Paris School of Economics, 2023 Leuven Summer Event, Toulouse School of Economics, NHH Bergen, Mannheim University, Queen Mary University, University Carlos III, the University of Geneva, University Cergy Paris, EARIE 2020, and the 2021 China Meeting of the Econometric Society for their feedback. Aru Bhardwaj, Miguel Blanco Cocho, Ian Hsieh, and Robert McDonald provided research assistance, with special thanks to Miguel and Ian. All remaining errors are my own. Early presentations of the work were titled "Competitive Award of Scarce Airport Slots: an Empirical Analysis". This project benefited from financial support by the *Sciences Po Scientific Advisory Board* and the *Banque de France - Sciences Po partnership.* Numerical computations were partly performed on the S-CAPAD/DANTE platform, IPGP, France, with technical support from Geneviève Moguilny.

1 Introduction

This paper investigates the welfare effects of implementing a market-based mechanism to allocate a scarce and highly valuable resource: airport slots. Airlines vie for take-off and landing rights twice yearly at over 200 of the world's busiest airports with limited capacity. A global airport slot allocation policy has been established to streamline this process across various airports. Key aspects of this policy include slots being distributed free of charge and airlines retaining their departure rights from the previous year as long as they cancel less than 20 percent of associated flights. The airport slot policy has garnered considerable attention due to its implications for airline competition, market efficiency, and potential windfall profits for slot-owning airlines. Amidst discussions on noise reduction, pollution control, and sustainable aviation, the need for a more efficient slot allocation policy is increasingly apparent. As flight movements come under scrutiny, comprehending the potential impacts of an alternative slot allocation system becomes vital for policymakers.

A well-designed auction mechanism assigns scarce resources to the firms that value them the most, typically leading to maximizing allocative efficiency as well. Market-based approaches are commonly employed for optimal assignments ranging from radio spectrum usage licenses to renewable energy production contracts. However, transitioning to auctions for airport slots may not necessarily enhance welfare. Firstly, a firm's valuation of a slot depends on the profits it can generate by operating flights in that slot, potentially leading to increased consumer costs. Secondly, since time slots are not tied to specific routes or flights, different airlines can utilize them for various markets, disrupting the alignment between allocative and productive efficiency and resulting in distributional implications. Moreover, the airline industry differs in that consumers derive some benefits from airlines having a larger share of airport slots due to associated convenience perks, a phenomenon widely documented.[1] It remains uncertain ex-ante whether these benefits outweigh the negative impacts stemming from reduced competition.

This paper presents the first welfare assessment of introducing auctions for airport slots. To accomplish this, I develop a flight-level structural model of demand and supply, adding significant detail to the existing literature. This model is estimated using a novel dataset of flights, which includes routes between two of the most slot-constrained airports, Paris Charles de Gaulle (CDG) and Paris Orly (ORY), and 119 global destinations accessible by direct flight during one flight season (summer 2018). A notable aspect of the model is its consideration of consumer preferences regarding flight departure times. Airline carriers apply higher equilibrium mark-ups for more popular departure times, which emerges as a key driver for the greater value of morning slots than afternoon slots. Furthermore, the model captures spillover effects across flights and markets, influenced by demand and supply factors. Estimating this model with a comprehensive global flight-level dataset adds to a deeper understanding of the airline industry, particularly as existing literature predominantly relies on US route-level data.

My empirical findings highlight the significance of the carrier's flight network. I observe diseconomies of scope in the number of destinations and economies of scale in the number of flights. Moreover, substantial scheduling efficiencies are evident. Specifically, a one standard deviation increase in the spacing of flights from the departure airport (by the same carrier and date and to any destination) leads to an 11.7 percent increase in costs, with other variables held constant at their sample means. On the demand side, consumers, particularly those in the business segment, place high value on flights offering better connection opportunities

[1] As documented by Borenstein (1989), Berry (1990), Brueckner et al. (1992), Berry and Jia (2010), Bontemps et al. (2021), Ciliberto et al. (2021), amongst others.

and are willing to pay a premium for them. These asymmetric network benefits are reflected in slot values. Consequently, transitioning to a market-based allocation of airport slots risks exacerbating concentration at the airport unless the mechanism is designed to favor small bidders or newcomers specifically.

In essence, a departure slot s grants the carrier the right to schedule one flight departing from the airport at the specified time and day of the week. The value of this slot to carrier a, denoted by $\nu(a, s)$, is defined as the incremental profits generated by the carrier when operating a flight in that slot. These profits are derived in market equilibrium and hinge on various factors such as the timing of the slot, the costs of and demand for the flight, and the role of the flight in the carrier's network. Additionally, the opportunity cost of winning the slot is not negligible. Slots hold greater value for a carrier when the slot use of one of its competitors is expected to reduce its profits significantly. This characteristic of the slot allocation issue leads to what is known as *slot hoarding* behavior, where carriers may even operate flights at a loss to retain their slot rights.[2] Considering all these factors, computing the full slot value would necessitate solving the market equilibrium numerous times, which is particularly demanding given the comprehensive demand and supply model with spillovers across flights and markets. I approximate two aspects of the slot value to speed up the analysis by a factor of about 240,000. Carriers are assumed to operate their best-performing available flight in the slot rather than the flight that would yield the most additional profit, and which flight they choose is considered public information. Section 2 provides further details about the proposed slot value characterization.

My approach involves aggregating all incremental profits for a flight based on the value of the departure slot at a single airport. This scenario assumes that only slots at one airport are auctioned while other airports continue their current process of awarding slots for free. Additionally, the slot value should be considered to apply to both the departure and reasonably close arrival slots at the origin airport. The estimated slot values are adjusted to apply across the 31-week flight season, creating a market for weekly slot pairs (WSPs) pertinent to this industry. This approach effectively captures the undeniable complementarity between arrival and departure slots at one airport and between *series* of weekly departure rights that facilitate network stability during the flight season. In addition, the nature of complementarity among WSPs is transparent and rooted in the model. An additional WSP reduces costs for all other flights operated on the same date to any destination by improving scheduling efficiencies and benefiting from economies of scale in flights. Furthermore, it increases demand non-linearly for all existing flights to the same destination on the same date, which can be interpreted as an expected delay effect. Numerical simulations of the model indicate that the complementarity across multiple WSPs is typically small since they are integrated into existing flight networks, limiting the impact of incremental cost savings or demand expansion. However, the complementarity can be significant for carriers with limited or no operations at the airport.

The median departure-arrival slot pair utilized for a weekly flight during a single flight season is estimated to hold a value of €579k. Air France, the largest carrier with over 50 percent of the slots in its global hub, tends to have higher slot values than other carriers for many slots. These discrepancies in slot values are partially attributed to variations in costs and demand. Air France benefits from a denser network and higher flight frequency to numerous destinations, giving it a competitive edge. Moreover, Air France operates some of the most lucrative flights from Paris to overseas destinations and is expected to expand its offerings in this regard. Comparing my findings to the scant information available regarding the prices paid for slots at London Heathrow Airport shows that my approach yields prices of a similar magnitude. While achieving a perfect

[2] See, e.g., De Wit and Burghouwt (2008), Fukui (2012), Miranda and Oliveira (2018), Fu et al. (2015), Kösters et al. (2023).

match is not the goal, obtaining similar prices validates the effectiveness of this approach that recovers auction values via their link with equilibrium profits.

To better understand the primary drivers of slot values, I conduct a regression of $\nu(a, s)$ on variables associated with the timing of the slot, the carrier's network, the flight, and competition at the route level. As anticipated, morning slots hold higher value; departing before 11 am increases the slot value by 25 percent. Slots at CDG carry greater value than those at ORY, albeit only on weekdays. All else being equal, slots are more valuable to carriers that presently possess them than to carriers competing for them, indicating the advantage of optimizing their flight network around the take-off and landing rights they currently hold. Moreover, slots command a significant premium for the dominant carrier at their global hub, even after accounting for flight connectivity and market competition. This advantage for the dominant carrier is equivalent in magnitude to utilizing the slot for a flight with either one standard deviation larger capacity or a destination one standard deviation further away. Given these variables, the first two moments of the idiosyncratic (residual) slot value distribution are higher for the dominant carrier than other carriers, suggesting that an asymmetric auction model would best characterize values in this market.

The estimated slot values are used to compute the welfare impacts of introducing an efficient mechanism in this market. The mechanism assigns each WSP to the carrier with the highest $\nu(a, s)$. I observe that the final ownership shares of the participating bidders are minimally affected by the auction; only easyJet experiences some loss of slots to other bidders, and apart from that the four bidders grow proportionally. Numerous slot trades occur, resulting in a substantial surplus for the median winning bidder. The absence of a secondary market for slots explains these findings; carriers typically cannot capitalize on opportunities to sell their least valuable slots and acquire new ones to enhance their network's value.

I use the model estimates to address the core indeterminacy in this environment as clearly identified in Borenstein (1988): efficient markets for licenses, which do not precisely determine prices, quantities, or even the applicable market, do not necessarily improve welfare or consumer surplus. I find that introducing a market for airport slots would be welfare increasing; re-assigning the 216 separate WSPs to the highest $\nu(a, s)$ firms is estimated to increase total welfare by approximately €96 million. Much of this is driven by €70 million higher consumer surplus. Due to the modeled spillovers through network effects coming through demand and supply factors, consumers in all markets are affected regardless of whether the product offering changes in those markets. The compensating variation differs substantially across markets, pointing to the policy's distributional impact, which leaves some consumers worse off. By contrast, the variable profits are estimated to increase for all participating firms as partly driven by the better utilization of the slots. Whether or not firms are better off overall depends on the exact auction mechanism (pinning down the price).

To gauge the scale of this market, I assume that the regulator would auction off only 10 percent of the slots each time to maintain continuity, with the winning bidder granted airport take-off and landing rights for 5 years. In this scenario, each auction involves approximately 900 slot series for the two airports in the study. Additionally, approximating the price as the second-highest value, akin to the Generalized Vickrey Auction, this is estimated to generate €2.2 billion in auction revenues. For context, these auction proceeds would be similar to those from the latest auction for 310 MHz of spectrum in the 3.4 - 3.8 GHz band conducted by the French communications regulator *Arcep*.

1.1 Literature

This paper aims to connect the two parallel pillars of research in Industrial Organization: equilibrium models of demand and supply for posted price markets (e.g., Berry (1992, 1994), Berry et al. (1995), Goldberg and Verboven (2001), Berry and Haile (2014)), and equilibrium models of bidding for auction markets (e.g., Paarsch (1992), Laffont et al. (1995), Guerre et al. (2000), Bajari and Hortaçsu (2003)). The auction literature leverages observed bids alongside game-theoretical equilibrium models of bidding to identify bidders' latent values (or their marginal costs), typically used to simulate the welfare impacts of alternative allocation mechanisms. I propose an alternative value-elicitation method, defining the item's value directly in terms of its profit-generating potential to the firm and estimating these profits with an equilibrium model of demand and supply in the industry. This framework is especially valuable for analyzing welfare effects of *introducing a market*, such as transitioning from a lottery or a first-come-first-serve allocation to an auction or matching mechanism or when revealed preference data from an existing market mechanism is unavailable. Doraszelski et al. (2017) is the only other paper I know performing such an ex-ante analysis, using an accounting model to estimate values for US broadband licenses before the auction. A benefit of my structural approach is that it enables counterfactual welfare analysis. Relatedly, Einav and Levin (2010) discuss parallels between demand estimation methods developed for auction and posted price markets.

The paper also contributes to the empirical air transportation literature (including Berry (1990), Berry (1992), Brueckner and Spiller (1994), Berry et al. (2006), Berry and Jia (2010), Ciliberto and Williams (2010), Aguirregabiria and Ho (2012), Gayle (2013), Ciliberto and Williams (2014), Lazarev (2013), Park (2020), Aryal et al. (2021), Bontemps et al. (2021), Williams (2022)). Noteworthy are Williams (2022) and Lazarev (2013), which, like this paper, depart from the typical disaggregated route-level analysis to consider the influence of flight characteristics (in their cases to study dynamic pricing behavior). Despite the industry's acknowledgment of the importance of airport slots dating back to Levine (1986), most literature abstracts from them. An exception is Park (2020), who evaluates slot divestitures in the American Airlines / US Airways merger, modeling demand for segments of routes and finding that the divestiture policy increased consumer surplus. Ciliberto and Williams (2010) provide empirical evidence of entry barriers that carrier investment in gates poses. Other studies, such as Borenstein (1989), Fukui (2012), Reitzes et al. (2015), Ribeiro et al. (2018), and Odoni (2020), focus on various aspects of airport slot policy without simulating equilibrium effects when the industry reacts to alternative policies. Grether et al. (1979), Rassenti et al. (1982), and related operations research papers examine the computational and operational feasibility of market-based airport slot allocation mechanisms. What sets this paper apart is its focus on the welfare implications of an efficient slot allocation mechanism, considering participating firms' equilibrium price and product strategies, while leaving the mechanism's design for future exploration.

The paper furthermore adds to the literature on auctions with after-markets. Efficient auctions typically bolster welfare by assigning items to the highest-value users (bidders). Still, welfare impacts become theoretically uncertain when these users compete downstream (Borenstein (1988), Janssen and Karamychev (2009), Hazlett and Muñoz (2009), Cramton et al. (2011), Gretschko et al. (2016), Kasberger (2020), Ershov and Salant (2022)). This literature predominantly focuses on broadband spectrum auctions and the associated telecom services market. Winning more spectrum licenses reduces the marginal cost of providing services across all markets, addressing the allocation of multiple marginal-cost-reducing licenses. While this applies to airport slot auctions due to estimated scheduling efficiencies and economies of density, additional demand-side network effects come into play. Moreover, as Borenstein (1988) points out, a crucial complicating factor is that the auctioned license does not specify the market or product it applies to — carriers can choose which flight to operate in the slot. Hence, both network effects and product choice are considered. A related line of research captures that firms first decide which markets to enter or products to offer before setting equilibrium prices (e.g., Eizenberg (2014), Wollmann (2018), Fan and Yang (2020), Park (2020), Bontemps et al. (2021)). Those papers follow the revealed preferences approach developed in Pakes (2010) and Pakes et al. (2015), necessitating that observed choices reflect (only) profit-based factors. When binding slot restrictions are introduced, the fixed entry costs (here, to operate a flight) can no longer be bounded. Instead, I approximate the optimal choice of product relying on the detailed accounting of demand and costs for a flight paired with operational constraints.

Finally, the proposed framework can be used to estimate bidder values for packages of items, contributing to a set of papers that perform empirical analysis of combinatorial auctions (Fox and Bajari (2013), Cantillon and Pesendorfer (2007), Kim et al. (2014), Kong (2021), Xiao and Yuan (2022), and Gentry et al. (2023)). Benefits of my approach are that it is independent of the auction format and that the estimated model parameters provide deep insight into the *drivers* of the complementarity.

2 The value of a slot

A departure (or arrival) slot gives the carrier the right to schedule one flight departing from (or arriving at) the airport at that specified time and day of the week. This flight generates a certain amount of profits for the operating carrier depending, in equilibrium, on the timing of the slot, the costs of and demand for this flight, the role of the flight in the carrier's network, and the flight offerings of competing carriers. I take all these factors into account. In this section, I define the value of an airport slot to an airline (section 2.1) and what this implies for the outcome of an auction for slots (section 2.2). In the next section, I propose a flight-level equilibrium model of demand and supply to estimate these slot values.

2.1 Defining slot values

First, some notation must be defined. Each slot ($s \in \mathcal{S}$) is associated with a departure airport, departure day of the week, and departure hour, the set of which is indicated by $\mathbb{D}_s$.[3] To indicate that flight j is operated in slot s, I indicate that its flight-level characteristics, $\mathbb{X}_j$, are contained in the set $\mathbb{D}_s$. To clarify that flight j departs from the departure airport associated with slot s, regardless of its departure day or time, I use $\mathbb{X}_j \in \mathbb{A}s$. Airlines are denoted by $a \in \mathcal{A}$ and markets by $t \in \mathcal{T}$. A flight is denoted by f and the flight that airline a offers in slot s is given by $f_a(s)$ (the determination of which is given below). The set of flights offered by all airlines across all markets is denoted by $\mathcal{F}$, and $\mathcal{F}^{s(a)}$ indicates the set of flights when airline a owns slot s.

Importantly, each slot can only be used for one flight. Hence, the difference in flight offerings when airline a or airline a' owns the slot consists of the two flights that they would operate in the slot, as in

$$\mathcal{F}^{s(a)} \setminus \mathcal{F}^{s(a')} = f_a(s) \text{ and } \mathcal{F}^{s(a')} \setminus \mathcal{F}^{s(a)} = f_{a'}(s). \quad (1)$$

[3]Slots are actually associated with a 5-minute time interval, but for simplicity and empirical feasibility, I consider hours. The model focuses on departure slots but can be applied to arrival slots as well.

Moreover, presuming that slots cannot be added to the congested system, the difference between the counterfactual sets of flights above and the baseline set of flights in the data are exactly two flights; the flight whose rights expired and that is taken out of the schedule ($f^X(s)$) and the flight that is added, such that $\forall a \in \mathcal{A}$

$$\mathcal{F}^{s(a)} \setminus \mathcal{F} = f_a(s) \text{ and } \mathcal{F} \setminus \mathcal{F}^{s(a)} = f^X(s). \tag{2}$$

Two initial simplifications regarding the choice of flight should be highlighted. Equations (1) and (2) implicitly consider that when one flight is added or removed from a carrier's network, the rest of the network does not change as a direct response. This is, to some extent, supported by the scarcity of slots and the complexity of the network. While airlines continuously work on improving their networks to meet demand best — for instance, all Air France flights to Chicago are moved to a new slot window in the summer season of 2019 (S19)—, the data also supports the lack of rigorous changes to the flight network when one flight is added or removed. This can be demonstrated by observing changes in Air France's flight network, which has the majority of departure slots in the data and is, therefore, in the best position to change its flight offerings. I focus on destinations where in S19 at most one more or fewer flights are offered than the same season the year before. In all six of these destinations, none of the added flights is offered in slot windows that weren't previously used to serve this market, and no slot windows are removed when reducing the offered flights by one. This means that new slots fit the existing *slot series* but for one additional or one fewer week. There is some reshuffling to the number of flights in each slot window, but this concerns less than 5 percent of the flights.[4]

Furthermore, when selecting which flight to offer at slot s, the airline is restricted to selecting a flight from their existing network ($\mathcal{F}_a$). This limits the risk of an exposure problem; only airlines who at baseline already operate from the airport are considered as competitors, and only flights to destinations that are already served at baseline can be added. It circumvents the idea that a critical mass of departure rights needs to be won by an (entrant) airline before it is worthwhile to offer flights from an airport.[5] Only considering the flight network of a carrier at baseline avoids having to model the entry decision of what would be the next best destination for an airline to expand its network to.[6]

The data justifies this approach. 92 percent of the flights newly offered in S19, by any carrier and to any destination reachable from Paris ORY and Paris CDG, go to a destination that a carrier flew to in S18. Some of the new destinations offered by carriers that were already present at the departure airport are operated with fewer flights than those offered by carriers new to Paris. This suggests that the fixed costs to offer flights to a new destination when already operating at the departure airport are lower than when setting up operations

[4] Table C1 in the online appendix documents these facts. The statistics indicate that network changes are sometimes made, but mostly flights are added in existing slot windows and a small share of flights in existing windows is updated between the two flight seasons. For example, the 4.17 for Calvi reflects that one flight has been added in S19 to the single slot time already used for 23 flights in S18. While several changes are made to the flight offerings to San Francisco (10 more flights are added to the Thursday 3 pm slot, 9 to the Thursday 3 pm slot, and 6 to the Wednesday 3 pm slot, and some flights are removed from the other slots), these changes constitute only a small share (0.23 percent) of all flights in those slots. Note also that S19 is one week shorter than S18, so we can tell from the table that Air France offers one weekly flight to Djibouti in both seasons. The change in the number of flights in that slot is 3.33 percent corresponding to 1/30 flights removed in the single slot window.

[5] Aguirregabiria and Ho (2012) estimate that (for a sample of the largest US routes) without non-stop connections in the two endpoint cities a carrier has to pay an entry cost of $536,000, whileacarrierwith53connectionsateitherendpointpaysonly$45,000 to start a new route between the two cities.

[6] Ciliberto et al. (2021) and Bontemps et al. (2021) model entry and pricing in airline markets, abstracting from binding slot constraints. Park (2020) considers slot constraints but, due to data limitations, can only generically consider them. Likewise, Ciliberto et al. (2021) and Bontemps et al. (2021) offer a simplified demand model, abstracting, for instance, from heterogeneity in consumer preferences that Berry and Jia (2010) find to be important, to estimate the two-stage entry and pricing model with spillovers across markets. My approach uses a rich second-stage model of demand and supply given the network of flights, combined with a simple entry choice, to obtain realistic and differentiated slot values accounting for network effects that spill over across flights and markets.

from scratch. All these forms of market entry are minor compared to expansions of the existing flight network.[7]

Airlines compete for a slot at auction, meaning that if they do not win, one of their competitors might. This matters if that competitor uses the slot to add a flight to a market that the airline itself also serves, as the competitor will steal some of its business. This feature of the slot allocation problem makes the opportunity cost of winning a slot at auction not zero, so that firms care about the winning bidder's identity even if they lose the auction themselves. The fact that the willingness to pay includes the value of deterring entry or expansion of competitors affects the efficiency of the market-based mechanism (as pointed out in Cramton et al. (2011)).

The choice of flight and dependence on competitors' flight choices complicates the estimation of slot values. To see why, consider a simplified case with only two airlines, a and a'. The value of slot s to airline a would be defined as the difference in total profits when a owns the slot and when a' owns the slot

$$\sum \mathbf{K}_a^*(\mathcal{F}^{s(a)}) - \sum \mathbf{K}_a^*(\mathcal{F}^{s(a')}) \tag{3}$$

subject to

$$f_a(s) = \underset{j \in \{\mathcal{F}_a | \mathbb{X}_j \in \mathbb{A}_s\} \cup \emptyset}{\arg\max} \left[\sum \mathbf{K}_a^*(\mathcal{F} \setminus f^X(s) \cup j) - \sum \mathbf{K}_a^*(\mathcal{F} \setminus f^X(s)) \right]$$
$$f^X(s) \text{ is known}$$

The vector $\mathbf{K}_a^*(\mathcal{F})$ stacks the equilibrium profits from all flights offered by airline a given $\mathcal{F}$, multiplying the vector of average flight-level markups $\mathbf{B}^a(\mathcal{F})$ with the equilibrium number of seats sold $\mathbf{Q}^*a(\mathcal{F})$, defined in section 3. Two complications arise when determining the full slot value defined in (3). The first is computational and relates to the choice of flight. Determining which flight adds the most value to an airline's flight network requires solving for the equilibrium in all affected markets for each of the flights (that depart from $\mathbb{A}_s$) in the networks of both airlines –as network effects and spillovers across flights and markets are considered relevant here.[8] That means that the market equilibrium needs to be computed $|\mathcal{F}| + 1$ times if all flights are offered from $\mathbb{A}_s$, for each of the affected markets, only to determine where the airlines would use their slots for. That is not a feasible journey to embark upon, with more than 120,000 products to consider.

The second complication relates to the role of private information. Equilibrium profits depend on flight-level demand and supply unobservables, typically assumed private information to the firm. Denoting the vector of ξ_j and ω_j $\forall j \in \mathcal{F}_a$ by ξ_a and ω_a, respectively, airlines can compute the exact values of

$$\mathbf{K}_a^*(\mathcal{F}) = \mathbf{K}_a^*(\mathcal{F}; \xi_a, \omega_a)$$

and thus determine the highest added profit flight $f_a(s)$. For their competitors' choice of highest added profit flight, they can only base themselves on expectations over cost and quality draws, deducting

$$\tilde{\mathbf{K}}_{a'}^*(\mathcal{F}) = \int_{\omega_{a'}} \int_{\xi_{a'}} \mathbf{K}_a^*(\mathcal{F}; \xi_{a'}, \omega_{a'}) dF_\xi dF_\omega.$$

[7] Table C2 in the online appendix documents these facts.

[8] Which markets are affected depends on the definition of the network effects. In our maintained specifications, the market equilibrium must be computed in all markets where a flight is added or subtracted from any network and in all markets where flights depart on the same date as a flight that has been added or subtracted.

That would also mean that the flight that a competitor will actually offer when obtaining the slot ($f_{a'}(s)$) is not usually the same as the flight that airline a expects it to choose ($\tilde{f}_{a'}(s)$). Computationally, it would require solving the market equilibrium $2\times|\mathcal{F}|+1$ times for each affected market; once to determine $K^*_a(\mathcal{F}\setminus f^X(s)\cup f_a(s))$ (to go into the slot value of firm a) and once to determine $K^*_{a'}(\mathcal{F}\setminus f^X(s)\cup \tilde{f}_a(s))$ (to go into the slot value of firm a'), and once when not adding a flight. Again, this is infeasible in an application with 120,000 products. To address this, I make the following assumption.

Assumption 1. *The flight operated by airline a in slot s is determined as*

$$f_a(s) = \underset{j\in\{\mathcal{F}_a|\mathbb{X}_j\in\mathbb{A}_s,\mathbb{C}_s\}}{\arg\max} K^*_j(\mathcal{F}) \tag{4}$$

and the flight whose departure rights expired is determined as

$$f^X(s) = \underset{j\in\{\mathcal{F}_a|\mathbb{X}_j\in\mathbb{D}_s\}}{\arg\min} K^*_j(\mathcal{F}) \tag{5}$$

and these flights are identified without error by all (competing) firms $a\in\mathcal{A}$. $\mathbb{C}_s$ denote additional constraints on the set of flights to be chosen from (detailed below), and $K^*_j(\mathcal{F})$ denotes the equilibrium profit of flight j given the baseline network of flights $\mathcal{F}$.

As such, the added flight $f_a(s)$ is taken to be the highest-profit flight to any destination currently served by the airline from the airport where s applies, but then taking off at the departure day and hour of s. The profits of each flight are computed in market equilibrium and denoted by $K^*_j(\mathcal{F})$. The dependence on the full set of flights indicates that equilibrium profits are not only a function of flights by the same or rival airlines within the market but also on flights by the same airline to other markets —through cost synergies. $\mathbb{C}_s$ in (4) denotes a set of additional constraints imposed when determining $f_a(s)$. These constraints guarantee that each flight can be replicated at most once, so that the frequency of flights to any destination can at most be doubled, and that the overnight status of the alternative flight has to match the overnight status of the flight that is currently served by the slot. Imposing $\mathbb{C}_s$ minimizes the scope for unrealistic out-of-sample predictions (such as airport limits on plane sizes, inoperability of certain slots for certain flight durations). The dropped flight, $f^X(s)$, is chosen as the one generating the least profit among all carrier flights (the slot owner) scheduled within the same departure slot as the one that expires.

The above determines the slot value for an airline with *one* competitor. Airlines obviously face multiple competitors on most routes, which is taken into account next. It is useful to mention that bidders only care about other firms in auctions without after-markets to the extent that more or more competitive firms will reduce their win probability for any given bid. Here, conditional on winning the slot, each competitor will affect the airline's expected profits across its flight network differentially. Previous studies document (De Wit and Burghouwt (2008), Fukui (2012), Miranda and Oliveira (2018), this US GAO report of 2012, or this research prepared for the EU Transportation Committee in 2016) or explain (Fu et al. (2015), Kösters et al. (2023)) airline behavior consistent with strategic *slot hoarding*, identified as the operation of flights with smaller aircraft and/or lower load factors and/or at higher flight frequencies than comparable flights in settings without slot constraints.

Kösters et al. (2023) include the following quote by the Lufthansa CEO Carsten Spohr, which is telling of the magnitude of this issue: *"We have to operate 18,000 additional, unnecessary flights during the winter,* *purely to secure our slots."* Worse than the operation of flights with sub-optimal load factors, airlines sometimes operate these so-called *rescue-* or *ghost flights* to safeguard their slot rights, despite substantial costs associated with having even (near-) empty aircrafts reach their destination.[9]

Such slot hoarding strategies require the airlines to value (substantially) the slot that does not fall into the hands of their competitors. This is accommodated in the slot value definition. Specifically, as formalized below, the slot value for airline a compares total profits when a uses the slot and total profits when the competitor whose slot use would decrease a's profits by the most.

Assumption 2. *The value of slot s to airline a is $\forall a \in \mathcal{A}$ and $\forall s \in \mathcal{S}$ given by*

$$\nu_{a,s} = \max_{a' \neq a \in \mathcal{A}} \left(\sum \mathbf{K}_a^*(\mathcal{F}^{s(a)}) - \sum \mathbf{K}_a^*(\mathcal{F}^{s(a')}) \right) \tag{6}$$

subject to the choice of $f_a(s)$ and $f^X(s)$ according to Assumption 1.

Two assertions can be made. First, the setting is one with *private values* because the realizations of competitors' cost and demand shocks do not affect $\nu_{a,s}$ and because $\nu_{a,s}$ is not affected by the $\nu_{a',s}$ of competing bidders (a'). [10] The private values paradigm adds to the tractability of this model in the first place because it allows these values to be estimated separately. Private values are also key for the counterfactual relying on the existence of an efficient mechanism. This relates to the second assertion that the setting is one with *allocative externalities* because K_a^* is a function of the full flight network. Jehiel and Moldovanu (2001) provide important results about conditions for the existence of Bayes-Nash incentive-compatible mechanisms in case of allocative externalities. They show that existence is guaranteed when values are private.

While this looks like a network formation or (route) entry game at first glance, there are important differences. A slot can only be used by one airline to operate one flight and markets are independent, so there is no best-response type behavior where if one airline enters a market the other airline should refrain from doing so. In other words, $f_a(s)$ is not a function of the choice of flight of competitors $f_{a'}(s)$. There can be one exception, as non-owners will know which rights expire ($f^X(s)$), so their choice of flight might depend on the flight that has been removed. This type of behavior would result in the slot value being subject to the outcome of a game between operators, and it is abstracted from by assuming that firms choose to replicate (drop) the flight that at baseline generates the highest (lowest) profits subject to constraints.

2.2 Linking slot values to auctions

Assumption 2 formalizes the slot value definition that is adopted throughout the paper. I assess the welfare effects of introducing a market for slots by allocating each slot to the airline that values it the most according to this definition and then comparing the counterfactual market equilibrium with the current one. This delivers results independent of the specific market design (as long as it is efficient). However, various auction rules are evaluated in the counterfactuals.[11] The results are realistic if the highest-value firm is correctly identified, even

[9]The public awareness of this issue peaked during the COVID-19 reduction in air travel (and the subsequent proliferation of ghost flights). Still, the issue predates COVID as documented by, for instance, this UK Competition and Markets Authority report of 2018. The quote from Carsten Spohr reflects the airline's strategic response to low demand during the COVID-19 period.

[10]This would not be the case if bidders either would not know the $f_{a'}(s)$ of their competitors or when bidders would incorporate the strength of those competitors in terms of win probability in $\nu(a, s)$.

[11]For example, the role of quantity caps and the presence of complementarities across slots are investigated. It is, however, not necessary to specify the auction mechanism further unless one is interested in the monetary outcomes of the mechanism (such as how much money will be raised by the auctioneer). While the results in this paper are informative for the mechanism's optimal design, a deeper analysis is left for future work.

if the true willingness to pay for the slot differs from (6).

To elaborate, the willingness to pay might differ because of certain unmodeled factors. First, the departure rights will be leased to airlines for multiple years to allow airlines to market and invest in routes properly (as in the Ball et al. (2007) proposal for LGA slot auctions), so the willingness to pay should reflect the discounted present value of the perpetuity of receiving $\nu_{a,s}$ for a set period of time. In simulations, I compute slot values based on 20 years at a 5 percent annual discount rate. I also multiply all slot values by the number of weeks in S18 (31 weeks) to reflect that carriers often operate on the same schedule for the entire season.[12] Second, $\nu_{a,s}$ ignores the fixed costs accrued by the extra flight. In fact, the marginal cost specification in (15) captures the average variable costs of offering the flight, including the average fuel cost per seat and other costs to operate the flight. Any additional fixed or variable costs incurred to offer the additional flight (purchasing a new aircraft, hiring more pilots, leasing an aircraft hangar, ...) are omitted.

All the factors described above are inconsequential for the results presented in this paper, as long as they do not vary substantially across airlines. In general, if we take $\mu_{a,s}$ to be the true willingness to pay of airline a for slot s, the results follow if $\mu_{a,s}$ is an affine transformation of the slot values defined in (6).[13] More formally, the counterfactual simulations represent true market outcomes if the following holds. The matrix ν of dimension $\big(|\mathcal{S}| \times |\mathcal{A}|\big)$ contains the values of all airlines for all slots as given in (6)

$$\nu = \begin{pmatrix} \nu_{1,1} & \cdots & \nu_{|\mathcal{A}|,1} \\ \vdots & \ddots & \vdots \\ \nu_{1,|\mathcal{S}|} & \cdots & \nu_{|\mathcal{A}|,|\mathcal{S}|} \end{pmatrix}$$

and the corresponding willingness to pay matrix of the same dimensions is denoted by μ. It is maintained that

$$\mu = \vec{a} \cdot \nu + \vec{b},$$

where $\vec{a}$ and $\vec{b}$ are $(|\mathcal{S}| \times 1)$ column vectors of factors outside the model that may vary across slots and that change the scale or the level of ν, respectively, but that are restricted to be the same for all airlines.

The counterfactual results rely on an efficient slot allocation mechanism that awards each slot to the airline that values it the most. The mechanism is implemented subject to various policy criteria, such as whether each airline can win at most a maximum share of all slots (i.e., a quantity cap), and should thus be understood to be constrained-efficient subject to such policy constraints collectively denoted by $\mathcal{P}$.

Assumption 3 summarises the above, guaranteeing that the winners are correctly identified.

Assumption 3. *Let $\nu^{\mathcal{P}}$ denote the matrix of slot values subject to the award mechanism's policy criteria and numerical implementation $\mathcal{P}$. The firm with the highest conditional slot value for slot s, indicated by $I\big(|\nu^{\mathcal{P}}|_s ==$*

[12] An ideal market-based mechanism would probably be designed to prioritize series of multiple consequetive slots, in line with the current allocation policy (see the World Airport Slot Guidelines and the discussion surrounding increasing the slot series length in Europe (e.g., this 2020 Airports Council International position paper). Table C2 shows that, in my data, one weekly slot is also the minimum frequency used to offer a new route when already operating at one endpoint; as Air France did for Bari and Ivalho in 2019, showing that auctioning longer slot series favors network expansion. The role of the series length is abstracted from in this paper by awarding all carriers weekly slots when winning.

[13] Note that it is not strictly necessary that ν is an affine transformation of the true willingness to pay μ. For simulating prices and other market outcomes after the auction, only the *identity* of the slot winner needs to be correctly identified, not the cardinal amount of their μ.

$\max(|\nu^{\mathcal{P}}|_s)\big)$, *is correctly identified as the winning bidder. Hence, the column vector of slot winners is given by*

$$\mathcal{I}^{\mathcal{P}} = \begin{pmatrix} \mathcal{I}_1^{\mathcal{P}} \\ \vdots \\ \mathcal{I}_{|\mathcal{S}|}^{\mathcal{P}} \end{pmatrix} = \begin{pmatrix} I\big(|\nu^{\mathcal{P}}|_1 == \max(|\nu^{\mathcal{P}}|_1)\big) \\ \vdots \\ I\big(|\nu^{\mathcal{P}}|_{|\mathcal{S}|} == \max(|\nu^{\mathcal{P}}|_{|\mathcal{S}|})\big) \end{pmatrix}$$

Implicit in the assumption is the existence of an efficient mechanism with market-clearing prices, which is not given in the presence of complementarities (e.g., Milgrom (2000), Cantillon and Pesendorfer (2007)). A specific mechanism with good properties that has been proposed (and tested experimentally) for auctioning slots at La Guardia is a (combinatorial) Clock Proxy Auction. Please refer to Section 8.1 for a discussion. The package bidding part is important for airlines to express synergies across slots. Four types of complementarities should be distinguished: 1) between arrival and departure rights at the same airport, 2) between departure and arrival rights for the same flight at two airports, 3) between departure (or arrival) rights in the same airport for multiple (consecutive) weeks in the flight season, and 4) between departure (or arrival) rights for multiple slots on the same date in the same airport.

The slot value definition in (6) loads all incremental profits for a flight on the value of the departure slot in one airport. This reflects the scenario where only slots in one airport are auctioned while the other airports stick to the current process of awarding slots for free. Moreover, the slot value should be considered to apply to the *pair* of departure and (reasonably close) arrival slots at the origin airport. As mentioned above, the values are scaled to apply to the 31 weeks of the season. Combined, this addresses complementarities 1-3 and results in a market for *weekly slot pairs* or WSPs, as is relevant for this industry. The auction could be implemented for multiple airports at the same time or for the weekly slot pairs at only one airport.[14] While complementarities 1-3 are undeniable, the fourth is interesting to investigate as it might be a key source of airline asymmetry. Therefore, the magnitude of synergies coming from demand (e.g., consumers valuing airport presence) or supply (e.g., scheduling efficiencies from sharing check-in desks) are studied in Section 7.2.

As in the empirical auction literature, one can interpret slot values as realizations of a random variable, $V \sim F_V$.[15] The slot value $\nu(a, s)$ is then a realization of V for carrier a in the auction for slot s, driven by a carrier-specific idiosyncratic value component ($\tilde{V}_a$) and a common value driver (the vector of observables $\mathbf{A}$)

$$V_a = \nu(a, s) = g(\mathbf{A}, \tilde{V}_a) \tag{7}$$

Variation in $\tilde{V}_a$ is generated by variation in ξ, ω, and by factors that affect $\mathbf{K}^*(\mathcal{F})$ but that are excluded from $\mathbf{A}$. Hence, in terms of order statistics, assumption 3 implies that the winning bidder in an auction with N bidders has a value of $V_{N:N}$; the maximum value of a sample of N draws from F_V. Section 7.1 reports estimated slot values and recovers both slot value determinants and $f_{V|\mathbf{A}}$ under additive separability.

[14] Implementing a slot auction in one airport works as long as its timing coincides with slot allocation periods in the current grandfathering policy so that bids can be made conditional on obtaining matching (arrival) slots in other airports.

[15] Bidder-specific subscripts are omitted from the notation for simplicity, but the results in Section 7.1 indicate that an asymmetric value model —distinguishing a large carrier in its global hub from other carriers— is more suitable in the empirical context.

3 Market Equilibrium conditional on Slot Allocation

This section presents a model for the airline industry suitable to capture slot values as defined in (6). Both departure-time preferences by consumers and airline-time-specific connection opportunities at the departure- and destination airports are key drivers of why some slots are more valuable than others. Including the importance of a route in the airline's flight network —in terms of cost and demand— accounts for hub effects at the origin and destination and facilitates the estimation of slot values on an airport-by-airport basis. The model takes as given both the aircraft fleet and the network of routes that are offered by the airlines, as these should be determined before the time when the slot auction is held. More details are provided below.

3.1 Equilibrium profits $\mathbf{K}^*(\mathcal{F})$

Multi-product firms compete in average flight fares (prices) for consumers within a market. As demand is assumed to be independent across markets, equilibrium prices can be obtained market-by-market despite the marginal cost being a function of the carrier's flight network ($C_{jt}(\mathcal{F}_a)$), as detailed below. Equilibrium prices maximize the firm's profits, solving $\forall(a,t)$

$$\mathbf{P}^*_{at} \equiv \{P^*_{jt}\}_{j\in\mathcal{F}_{at}} \equiv \arg\max_{\mathbf{P}_{at}} \Big(\sum_{j\in\mathcal{F}_{at}} \big(P_{jt} - C_{jt}(\mathcal{F}_a)\big) Q_{jt}(\mathcal{F}_t, \mathbf{P}_t) \Big). \tag{8}$$

$Q_{jt}(\mathcal{F}_t, \mathbf{P}_t)$ denotes the quantity demanded given 1) the vector of prices $\mathbf{P}_t$ for all flights in the market and 2) the set of flights $\mathcal{F}_t$ (as the number of flights to a destination is allowed to affect consumer demand). It equals the size of market t times the equilibrium market share for product j ($S^*_{jt}(\mathcal{F}_t, \mathbf{P}_t)$ defined in (13)). From this, standard Bertrand-Nash first order conditions for optimal prices are derived, requiring that $\forall j \in \mathcal{F}_{at}$

$$\sum_{h\in\mathcal{F}_{at}} (P^*_{ht} - C_{ht}(\mathcal{F}_a)) \frac{\partial Q^*_{ht}(\mathcal{F}_t, \mathbf{P}^*_t)}{\partial P^*_{jt}} + Q_{jt}(\mathcal{F}_t, \mathbf{P}^*_t) = 0. \tag{9}$$

The equilibrium profits are derived without observing the marginal costs by subtracting the equilibrium mark-up from observed (equilibrium) prices. Rewriting (9), the mark-ups are given $\forall(j,t)$ by

$$B^*_{jt}(\mathcal{F}) = -\Big(Q_{jt}(\mathcal{F}_t, \mathbf{P}^*_t) + \sum_{h\neq j\in\mathcal{F}_{at}} (P^*_{ht} - C_{ht}(\mathcal{F}_a)) \frac{\partial Q_{ht}(\mathcal{F}_t, \mathbf{P}^*_t)}{\partial P_{jt}} \Big) \Big(\frac{\partial Q_{jt}(\mathcal{F}_t, \mathbf{P}^*_t)}{\partial P^*_{jt}} \Big)^{-1}. \tag{10}$$

$B^*_{jt}(\mathcal{F})$ equates to the standard multi-product Bertrand-Nash equilibrium mark-up of product j in market t. However, in this model, it depends on the network of all flights in the market ($\mathcal{F}_t$) and the network of own flights to any destination ($\mathcal{F}_a$), which for notational simplicity is denoted by the full set of flights $\mathcal{F}$. The equilibrium profit of a flight is

$$K^*_{jt}(\mathcal{F}) = B^*_{jt}(\mathcal{F}) Q_{jt}(\mathcal{F}_t, \mathbf{P}^*_t). \tag{11}$$

This also defines $\mathbf{K}^*_a(\mathcal{F})$, which stacks the K^*_{jt} $\forall(j,t) \in \mathcal{F}_a$, and which is used to determine slot values to firm a in (6).

3.2 Ticket demand

The demand model is based on Berry et al. (2006) and Berry and Jia (2010). Two latent consumer types vary in their preferences for product characteristics, containing at least the price of the flight. The two types $r \in \{1, 2\}$ are interpretable as consumers traveling for either business or leisure purposes. When choosing among products, purchasing a different flight and not flying altogether are categorically different, justifying a nested logit specification with one nest for all flights in a market and one nest for the outside option ($j = 0$). As such, the indirect utility for consumer i of type r for consuming flight j by carrier a to market t is given by

$$U_{ijt}(\mathcal{F}_{at}) = \alpha_r P_{jt} + X_{jt}(\mathcal{F}_{at})'\beta_r + \epsilon_{ijt}(\kappa), \tag{12}$$

and $U_{i0t} = \kappa\zeta_{i0t}$ for the outside option of not flying. The composite error term $\epsilon_{ijt}(\kappa)$ equals $\xi_{jt} + \nu_{it}(\kappa) + \kappa\zeta_{ijt}$; ξ_{jt} captures the flight-level scalar that is unobserved to the econometrician but (potentially) observed by firms and consumers, ν_{it} captures consumers' preference for flying over not flying to market t, ζ_{ijt} captures consumer i's idiosyncratic utility for flight j to market t. Parameter $\kappa \in (0, 1]$ governs how closely substitutable flights are relative to the outside option. If $\kappa = 1$, the model reduces to a simple logit model where substitution to the outside option is as likely as substitution to another flight within the same market. It is maintained that $\mathbb{E}[X_{jt}(\mathcal{F}_{at})'\epsilon_{ijt}(\kappa)] = 0$ but $\mathbb{E}[P'_{jt}\epsilon_{ijt}(\kappa)] \neq 0$ due to equilibrium prices being (potentially) a function of unobserved flight attributes. Furthermore, ζ_{ijt} is assumed to be Gumbell-distributed, and ν_{it} is drawn from the distribution shown by Cardell (1997) to render $\epsilon_{ijt}(\kappa)$ Gumbell-distributed as well, to generate the familiar closed-form nested logit purchase probabilities.

In particular, the market share of product j in market t is given by

$$S_{jt}(\mathcal{F}_t; \xi_t, \theta^D) = \gamma_1 D_{j1} \frac{\left(\sum_{k=1}^{J} D_{k1}\right)^{\kappa-1}}{\left(1 + \sum_{k=1}^{J} D_{k1}\right)^{\kappa}} + (1 - \gamma_1) D_{j2} \frac{\left(\sum_{k=1}^{J} D_{k2}\right)^{\kappa-1}}{\left(1 + \sum_{k=1}^{J} D_{k2}\right)^{\kappa}}, \tag{13}$$

where

$$D_{jr} = \exp\left((\alpha_r P_{jt} + X_{jt}(\mathcal{F}_{at})'\beta_r + \xi_{jt})/\kappa\right).$$

γ_1 is the share of the population that is of type 1 and $\theta^D = (\kappa, \gamma_1, \alpha_1, \alpha_2, \beta)$ is a vector of all demand parameters. The market share comprises the share of consumers of type r that purchase any ticket and the share of those consumers that choose flight j conditional on flying. The notation makes explicit that S_{jt} depends on the vector of all product-level unobservables in the market (ξ_t) as well as the set of flights offered in the market ($\mathcal{F}_t$).

The demand model is specified to recover slot-related consumer preferences. Consumers are allowed to prefer one departure airport over another and to have preferences for certain departure days and/or times. Hub effects are accounted for in demand through the dependence of $X_{jt}(\mathcal{F}_a)$ on the airline's flight network. The importance of a carrier's airport presence for demand is well-documented in the literature (Borenstein (1989), Berry (1990), Brueckner et al. (1992), Berry and Jia (2010), Ciliberto et al. (2021)) and is accounted for by including the number of destinations serviced from the departure airport. Moreover, consumers likely prefer to fly with a carrier who has more frequent flights to the destination as this reduces wait time when a flight gets cancelled. Together with the variables capturing the potential connections at the origin and destination airports described above, these can be considered demand-side hub effects. Including these in the demand model makes it unclear whether consumers are better off when taking a slot from the dominant carrier and giving it to a competitor in the same market, even when prices decrease. The estimated demand parameters will illuminate a fundamental trade-off between higher prices and more convenience associated with concentration in this industry. Consumer utility also depends on airline (group) fixed effects and the distance to the destination to proxy for the feasibility of other transport modes such as a high-speed train.

Overall, this constitutes a rich model of consumer preferences, especially in relation to a carrier's network of flights and the role that a certain departure slot could have in it.

The flight-level equilibrium model differs from the more standard itinerary-level analysis of the industry, and the distinction is important to make here. An *itinerary* is a set of flights between origin and destination (and, for a round trip, back to the origin) for which a passenger buys a ticket. Hence, while itineraries have connections, flights usually do not. In my Global Schedules Data, only 128 out of 152,531 flights (0.08 percent) that departed from CDG and none that departed from ORY in S18 were connecting flights. By contrast, *Aéroports de Paris* documents that 21.7 percent of the passengers departing from its airports (CDG and ORY) were connecting from another flight that year.[16] The large discrepancy between these two statistics originates from the difference between an *itinerary* as viewed from a consumer's perspective and a *flight* as scheduled by an airline carrier. The flight-level model is developed to capture that an operating carrier must have arrival and departure rights for each slot-constrained airport on the itinerary. Moreover, the carrier's willingness to pay for those slots is derived from the additional profits generated (across the carrier's network) with the single flight that can be scheduled in a slot.[18]

A non-negligible portion of demand comes from passengers who are on multi-legged itineraries. I address this as follows. First, the market size definition accounts for connecting passengers (see section 4.3). Second, the utility function allows for a demand premium for flights with better connection possibilities to other flights by the same carrier or codesharing partners, both at the origin and the destination airport. Specifically, the demand specification includes the number of other flights by the same carrier or a codesharing partner that arrive in (depart from) the departure airport (arrival airport) within a time window of 1.5-6 hours before the flight takes off (after the flight arrives). These variables reflect the convenience of the flight in relation to other flights on a (potential and unobserved) itinerary, as also typically accounted for in itinerary-level analyses (e.g., the number of connections in Berry and Jia (2010), the ratio of itinerary distance over non-stop distance in Gayle (2013)). Third, while some consumers face prices for multi-legged itineraries, my price data contains the average return fares for all direct flights between CDG and ORY and the destination airports. It is therefore maintained that the (unobserved) fare for the full itinerary consists of the sum of ticket fares for all of its legs — which Williams (2022) finds to hold for most observations in his sample of daily ticket prices for trips in US monopoly markets.[19] My quantity data already details the number of passengers on each flight, irrespective of

[16] Source: press release *Aéroports de Paris* January 15 2019, available here: http://hugin.info/145257/R/2231408/877211.pdf, last accessed 05-05-2022. The flight statistics are based on the Cirium global schedules file. Overall, it contains the universe of 32,342,819 flights for the 2018 summer season, of which 26,800,445 are scheduled passenger flights relevant for this study.[17] 86.6 percent of all scheduled passenger flights globally is indicated to be direct, 10.5 percent has one stop, and the remainder has two or more stops. The 128 flights with connections at CDG are for only 6 routes, documented in table C7.

[18] Figure C4 illustrates flight-level analysis in relation to the value of a slot. In the figure, the value of slot is derived from the incremental profits to the airline when offering flight CDG-MPX (assuming this flight is selected as $f_a(s)$), irrespective of the itinerary that the passengers on that flight are on. The equilibrium profits are based on the number of occupied seats on the plane (quantity), the average ticket fare of the flight, flight-level marginal costs, and spillovers of costs and demand to other flights by the same carrier.

[19] A similar rule applies to studies that use the DB1B dataset. That data includes the price paid for the full itinerary, and in the case of round trips, this price is divided by two for each of the two ways for a closed-jaw ticket (with the same or no stops in either direction). Moreover, the fare is typically assumed to be proportional to the miles in each direction otherwise. This procedure is explained on the DB1A/DB1B page of the NBER data archive: https://www.nber.org/research/data/department-transportation-db1adb1b, last accessed 05-05-2022.

where they originated from and whether they took another flight from the destination. On the supply side, the marginal costs to operate a flight are taken to be independent of the number of flights on a passenger's itinerary as it is unlikely that the sequence of flights that a passenger takes influences the cost to operate those flights.[20]

3.3 Marginal costs

The marginal cost per seat of flight j to market t is obtained from (10) as

$$C_{jt}(\mathcal{F}_a) = P_{jt} - B^*_{jt}(\mathcal{F}; \hat{\theta}^D) \tag{14}$$

where $B^*_{jt}(\mathcal{F}; \hat{\theta}^D)$ indicates the equilibrium mark-up given the estimated demand parameters. The costs are further decomposed into

$$C_{jt}(\mathcal{F}_a) = W_{jt}(\mathcal{F}_a)'\theta^S + \omega_{jt}, \tag{15}$$

where W_{jt} contains observed aircraft- and distance-specific fuel expenditure, other aircraft specifications, and variables depending on the carrier's own flight network —capturing scheduling efficiencies—, and ω_{jt} is an additive idiosyncratic cost term assumed to satisfy $W_{jt}(\mathcal{F}_a) \perp \omega_{jt}$. C_{jt} is not required to be strictly positive. In fact, taking economies of density and airlines' hub-and-spoke strategies seriously, C_{jt} is more accurately interpretable as the average variable cost per seat on flight j *net of any benefits* of transporting passengers to destination airport t. For instance, we might observe negative costs when fares are substantially discounted on feeder flights to a hub.

Fuel expenditures comprise a large share of the variable costs of operating commercial aircraft. The prior literature typically approximates this cost by flight distance (e.g., Berry and Jia (2010), Gayle (2013), Bontemps et al. (2021)). Observing the aircraft type in my data allows for a richer approximation. The average fuel consumption per seat is approximated by an aircraft-specific quadratic function in the distance to destination with aircraft-specific fuel efficiency coefficients estimated in Seymour et al. (2020), which they show mimics more complex physics-based fuel burn performance models used in the industry.[21] Fuel consumption per seat is multiplied by the Kerosine-type Jet Fuel spot price one month before the departure date to (imperfectly) account for input price variation over time.

In addition, the carriers' flight networks naturally have a bearing on costs, mainly due to economies of density associated with a hub-and-spoke network compared to a point-to-point network. To capture differences in (counterfactual) profits across airlines, it is important to quantify these network-based cost differences and what drives them. The big four airlines in my data use different network strategies, as documented in Figure 1. Air France (The Air France-KLM Group) and Lufthansa have clear hub-and-spoke flight networks. Through the lens of this study, Lufthansa's hubs are at destinations (Frankfurt and Munich), whereas Air France has a big hub at the origin (Paris) and some destinations (Amsterdam, for instance). The flight networks of the next two biggest carriers, easyJet and Vueling Airlines, reveal more of a decentralized network strategy, although

[20] This is slightly different from the approach in Park (2020), where an itinerary-specific unobserved error term is added to deterministic marginal cost components that are additive over multiple legs of the itinerary.

[21] Such so-called "high fidelity" physics-based models are used to simulate aircraft performance's geometric, kinematic, and kinetic properties. They are complex and require a wealth of data such as cruising altitude, lift coefficient, climbing speed, and flight envelope (Seymour et al. (2020)). See for instance https://www.eurocontrol.int/model/bada for EUROCONTROL's aircraft performance model.

(a) AirFranceKLM

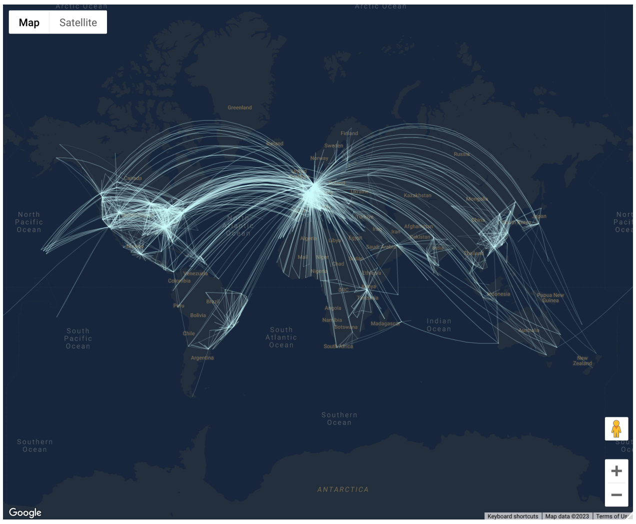

(b) Lufthansa

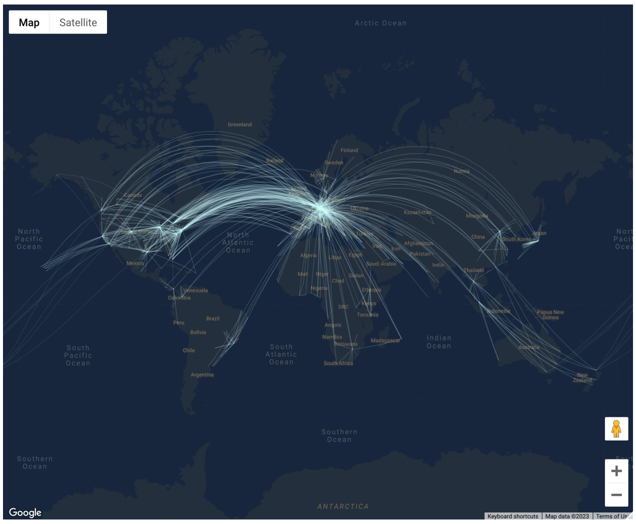

(c) AirFranceKLM (Europe)

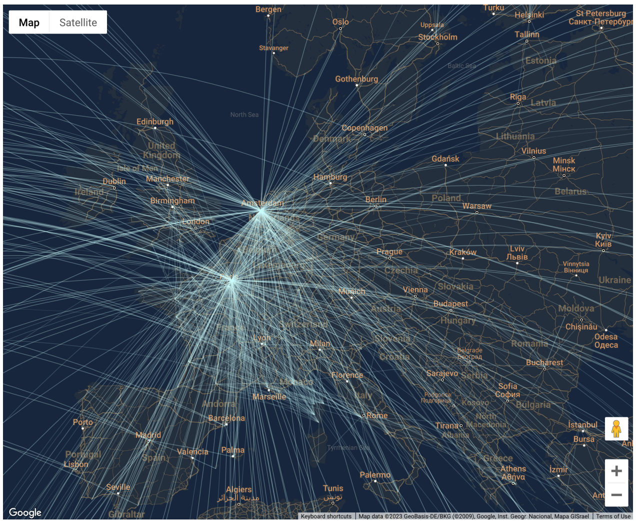

(d) Lufthansa (Europe)

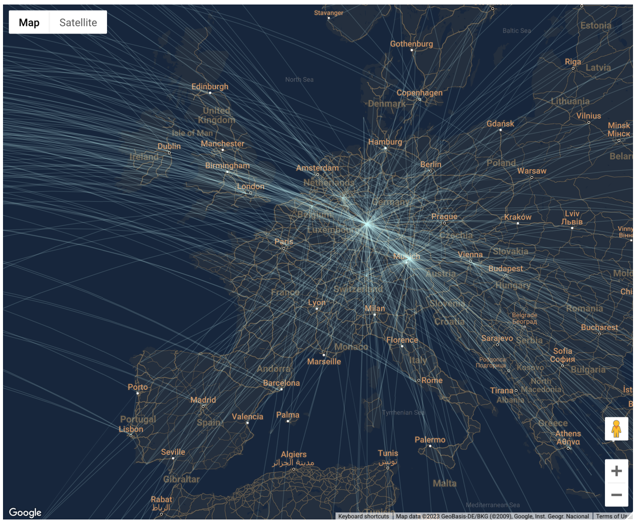

(e) easyJet

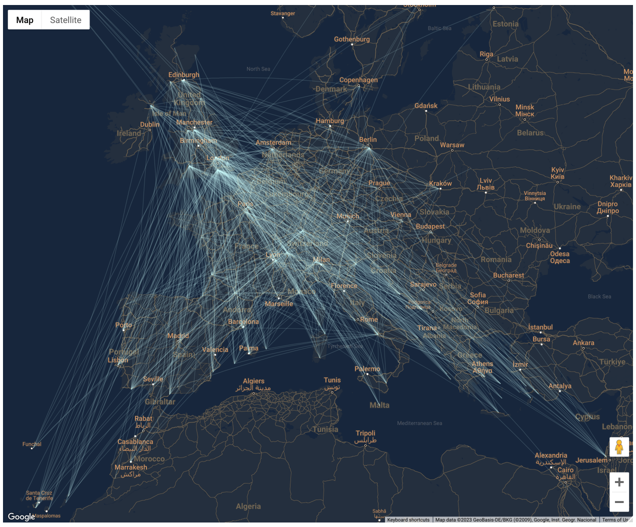

(f) Vueling Airlines

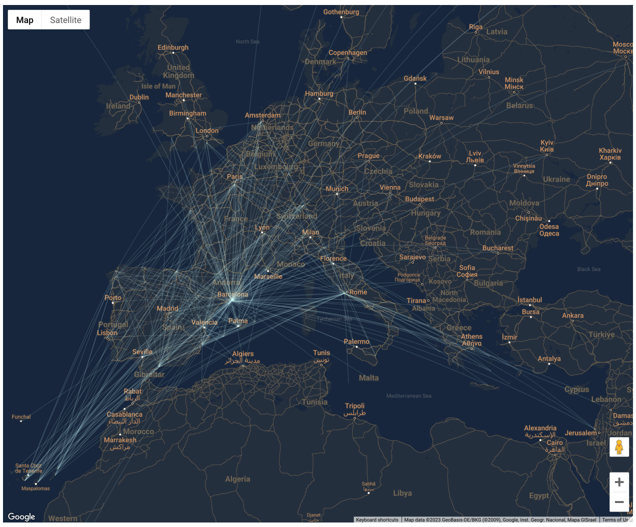

Figure 1: Global Flight Networks of largest 4 carriers at Paris CDG and ORY.
Based on the Global Schedules File for S18. Flights where the carrier is the marketing carrier in a codesharing agreement are included.

with disproportionally many departures from airports in the UK and Spain, respectively.[22]

Note that in this flight-level model, economies of density are to a large extent already captured by the flight's load factor.[23] Nonetheless, to sustain a high load factor, transporting passengers from or toward hubs is beneficial, for instance, by the cost estimates for US flights in 2006 reported in Berry and Jia (2010). For our purposes, it is relevant to quantify whether and by how much the slot value increases when airlines can use that slot for a flight with at least one leg in a hub. As such, the flight-level marginal cost function includes as explanatory variables the number of destinations the carrier serves from the departure airport and the number it serves from the arrival airport.

Beyond the load factor or the number of flights on a route, the timing of those flights might also be relevant for costs. Flights departing in close proximity to each other can have check-in personnel take care of both flights, for instance. As such, the average spacing in hours of adjacent departures (of all flights by the same carrier departing on the same date to any destination) is included in the cost function. These variables are computed at the origin and destination, and are inspired by the departure time differentiation measures in Borenstein and Netz (1999).

As a result, the model contains both cost synergies and market spillovers. Consider an airline that receives a slot and adds a flight to a certain destination. If consumers value the frequency of flights to a certain destination, their demand for *all* flights by this carrier to that destination increases, ceteris paribus. Moreover, the airline potentially enjoys cost savings as the number of flights to any destination decreases, and this efficiency gain would lower the marginal cost of all flights to any destination. In addition, he might obtain a cost reduction for all flights departing on the same date. These network effects will be estimated and require the summation of profit changes overall markets when computing the value of a slot in (6).

Finally, I follow most of the literature and omit a fixed cost (here, to operate a flight). This abstracts from the decision to cancel a flight when the realized demand is insufficient to cover fixed costs. In the current setting, the "use-it-or-lose-it" rule that applies at both slot-constrained departure airports in the sample makes it particularly unattractive for a carrier to cancel a flight, justifying the abstraction. Moreover, fixed flight costs are only relevant in this study insofar as they vary across flights and/or airlines (as explained in section 2.2), above and beyond factors already included in the relatively detailed cost specification given in Table C9.

4 Data description

This section describes the data that is used to estimate the model.

4.1 Data sources

Two main data sources are used for the analysis. The first is data from French civil aviation regulator DGAC containing a sample of ticket prices and number seats occupied for all return flights departing from CDG or ORY in S18. The administration uses this data to construct the travel component of the Consumer Price

[22] In the context of severe (slot) constraints on expanding at airports, it is likely not by choice that these two low-cost carriers did not manage to replicate the degree of network centralization of the *flag carriers* — often privatized national carriers that still enjoy privileges. Airlines' historic rights to slots are granted based on the flight network when the airport reaches a certain level of congestion. As these rights carry forward, it is difficult for newcomers to accumulate a substantial share of slots.

[23] Relatedly, the number of flights between origin and destination does not need to be included on the cost side or instrumented for on the demand side (as in, e.g., Berry and Jia (2010)), with the unit of observation being a flight rather than a composite product of multiple flights capturing all demand between origin and destination in a period.

Index. This unique dataset allows for studying the slot allocation policy with the proposed methodology, as it contains flight-level identifiers, including its departure time and date. Such information is not available in the DOT DB1B 10% random sample of US tickets normally used in empirical airline papers. Like the few other papers that have collected flight-level data to address questions that cannot be answered with the DB1B (i.e., Lazarev (2013), Williams (2022), and Aryal et al. (2021)), the price data is based on offered fares.[24] More specifically, the price variable is constructed as the lowest offered fare weighted based on passenger surveys to reflect actual purchase prices for passengers flying through CDG and ORY airports in S18. The survey weights are a particularly nice feature as it makes the survey-weighted lowest offered fare a close proxy of actual prices paid by the the relevant population. The DGAC quantity data compares well to the widely used data from the US Department of Transportation.[25].

Despite this dataset's high quality, the DGAC's algorithm to obtain prices does not regard the flight's departure time. As a result, the price is missing for a large share of flights for which a lower-priced alternative (with otherwise identical features) took off at a different time that day. Online Appendix A details the fare sampling process and motivates the method to correct for fares not missing at random. To summarise, when ξ is a proper flight-level unobservable, constant across multiple offers (each time DGAC samples the offered fares) and independent of any offer-specific unobserved heterogeneity, the data is missing at random conditional on observables and ξ. As such, imputation delivers a consistent and asymptotically normal estimator (Dagenais (1973), Gourieroux and Monfort (1981)) when capturing ξ with observables such as S_{jt}.[26] Random Forest (in particular, the `ranger` implementation in `R`) performs best across a range of machine learning- and regression models and is adopted to impute the missing fares. The $R^2 = 0.89$ in the test sample that is left out from the data on which `ranger` searches for the best partitioning (to guard against over-fitting). Online Appendix B provides further details.

The second data source is the Global Direct Schedules File purchased from Cirium. It contains all scheduled flights globally for the period January 2018-October 2019. This data is particularly useful for a rich analysis of how flight networks affect slot values, as it is not limited to Paris-destination airport pairs.[27] Additional data is obtained from publicly available sources; the *US Energy Information Agency* for weekly jet fuel prices, *Simplemaps* for the 2019 population count of cities, and *Google Maps* for their coordinates.[28] Finally, aircraft-specific fuel consumption performance parameters are taken from Seymour et al. (2020).

4.2 Estimation sample and summary statistics

The estimation sample contains flights departing from Paris Charles-de-Gaulle (CDG) and Orly (ORY) airports in the S18 (between 26 March and November 3rd, 2018). While the quantity data is available for all 300 destinations reachable from CDG and ORY by commercial flight, the DGAC collects prices only for a subset

[24] See also Ivaldi et al. (2022), who describe the limitations of both the DB1B and of flight-level data based on offered fares.

[25] The US DOT does not publish prices for flights departing or originating outside of the US, but the DOT T-100 dataset reports the number of tickets sold for transatlantic flights. Passenger shares by US destination (for destinations that are in both samples) are very similar to those in the DGAC dataset (see Table C3 in the online appendix)

[26] Table B3 shows that the market share conditional (*prodshare*) and unconditional (*marketshare*) on purchasing, and the share of the outside good (*outsideshare*) are among the 50 most important predictors of P_{jt}. Market shares are suitable for this prediction problem as a proxy for ξ conditional on other observables, but it would be misguided to attach a causal interpretation to results from a regression of prices on market shares (Miller et al. (2022)).

[27] Statistics from this dataset are averaged over different schedule observations weighted by the number of seats in the aircraft if needed to account for a carrier's schedule variation within the season.

[28] The US EIA publishes historic U.S. Gulf Coast Kerosene-Type Jet Fuel Spot Prices (dollars per gallon) here: https://www.eia.gov/dnav/pet/hist/, the Simplemaps data can be downloaded here: https://simplemaps.com/data/world-cities. Ggmaps in R is used to obtain city coordinates to compute distances.

of those destinations. The 119 destinations in the estimation sample are more important in the sense of the number of carriers and the frequency of flights that go there. The most popular destinations for ticket sales are Nice, Toulouse, Barcelona, Madrid, Montreal, Lisbon, and New York. All cities are kept in the sample regardless of their population size, as small cities by that definition (such as Toulouse) are also important by demand volume for the relevant population. The sample is restricted to the 57 operating carriers for which the DGAC collects fares. These carriers are more relevant in the sense of serving more destinations and offering more flights to any destination. In the few instances where multiple flights by the same carrier depart to the same destination in the same hour on the same date, the later flights are dropped from the sample to retain the model's flight-level nature —also when specifying departure time preferences as hourly fixed effects. The final estimation sample contains 120,0039 flights (observations).

On average, a flight has 176 passengers. The average load factor (the number of occupied seats divided by the aircraft's capacity) is 0.85. Only five percent of flights are full. By comparison, in the sample of US monopoly markets of Williams (2022) the average load factor is 89 percent but 15.7 percent of flights sell out. So-called ghost flights are not important in the data.[29] Only about 0.2 percent of flights takes off with less than 25 passengers and about 0.6 percent of flights takes off with less than 25 percent of its seats occupied.

The average price for a flight is €302, with a relatively large standard deviation that reflects the mix of low-cost short-distance flights and more expensive long-haul flights. The four quadrants break down this variation by the departure airport and whether the flight is overnight. There is a large variation in fares within and across departure times and more so in CDG than at ORY. ORY is a smaller airport closer to the city, with more short-haul flights than CDG. The average flight distance in the estimation sample is 1629 km (1012 miles) at CDG and 972 km (604 miles) at ORY. 48 of the 111 CDG destinations are also accessible from ORY, and 8 can only be reached from ORY. Both are among the most slot-constrained international airports in the world.

Among the 57 carriers, Air France stands out as the dominant carrier in both airports.[30] Air France was the national flag carrier before the merger with KLM in 2003, is headquartered in Paris, and operates 53 percent of all flights in the data. The online appendix contains additional details supporting the statistics referenced in this section.

4.3 Product and market definition

A product is a direct flight between origin and destination, consisting of the operating carrier, the departure time and date, and the average fare. The estimation sample has only one airport per destination city. Markets are therefore all cities that can be reached by non-stop flight from Paris, conform the directional city-pair market definition but limited to departures from Paris in line with the scope of the analysis. The market definition also contains the departure month, based on the assertion that substitution between flights in different departure months is limited. It is valid when their travel demand is associated with a fixed period (say, due to a conference or a school vacation). This reduces the average number of products per market from 952 to 119, and it does not affect the competitive landscape regarding the number of carriers per market. The number of carriers per

[29] A ghost flight is a flight that takes off with none or very few passengers, usually motivated by the 80 percent "use-it-or-lose-it" requirement of the current slot allocation process. Flights may also take off (near-) empty for other reasons, for instance, if the aircraft is needed to operate a flight from the destination. A frequent occurrence of ghost flights would motive departing from the standard Bayes-Nash profit maximizing equilibrium.

[30] In what follows; all three carriers in the Air France Group are collectively referred to as Air France, and they are considered to maximize the group's profits. Carriers of the Air France Group included in the estimation sample are Air France, KLM, and Transavia France.

market is hardly affected because carriers typically maintain a regular flight schedule where they offer the same flight throughout all months of the flight season.

Adjusting the usual market size definition accounts for connecting passengers. The market size is made up of the geometric mean of the population in the destination city and the population in Paris weighted by the share of non-stop passengers in Paris reported by *Aéroports de Paris* (0.873) plus the average geometric population mean across all markets weighed by the share of connecting passengers (0.217). This reflects that a non-negligible share of demand for flights departing from Paris is by passengers flying from or to cities beyond the origin or destination of the flight. Assuming finally that people need to travel about twice per year, the market size equals

$$M_t = \frac{1}{6}\left(0.783\sqrt{(pop_{Paris} \times pop_t)} + 0.217\frac{1}{T}\sum_t \sqrt{(pop_{Paris} \times pop_t)},\right)$$

where T denotes the total number of markets. Population size (pop_t) is constant across months for the same destination.

5 Estimation method

This section details how the model is estimated, what variation in the data-generating process identifies the parameters, which instruments are used, how the estimator is implemented numerically, and on what basis the final model specification is selected.

5.1 Estimation and Implementation

The demand model is a random coefficients nested logit (RCNL) model (previously adopted in, e.g., Brenkers and Verboven (2006), Miravete et al. (2018), and Miller and Weinberg (2017)) but with discrete consumer types (travelling for business or leisure) as in Berry and Jia (2010). In this approach, a generalized method of moments (GMM) estimator is derived from the exclusion restrictions of the instruments for the price (Z^D, described below), namely that

$$\mathbb{E}[\xi(\theta^D)' Z^D] = 0. \tag{16}$$

In this equation, $\xi(\theta^D)$ is the vector of flight-level structural error terms in consumer utility (12) that equate the observed market shares (the vector $\hat{\mathbf{s}}_t$) to the market shares predicted by the model at the true values of the parameters as defined in (13), e.g. the implicit function that, in all markets, sets

$$\hat{\mathbf{s}}_t = [S_{0t}(\mathcal{F}_t; \xi_t, \theta^D), \ldots, S_{Jt}(\mathcal{F}_t; \xi_t, \theta^D)]'. \tag{17}$$

I apply the standard fixed point algorithm of Berry et al. (1995) to solve the implicit function, using the modified contraction mapping that Grigolon and Verboven (2014) show accounts for the nesting structure of the error term. I find that it is important to include destination fixed effects in the demand model. In specifications without destination fixed effects, even though the estimates still reveal a preference for morning flights, the coefficients on the departure hour dummies are driven toward zero. This is because the availability and price of alternatives vary substantially by destination. Not controlling for destination effects mixes the utility of flying to the destination airport with the utility of departing at the time of the available flights. More details are given in section 5.3.

The destination fixed effects are concentrated out in the estimation algorithm, meaning they are estimated for each set of candidate parameters $\tilde{\theta}^D$ that the optimizer computes the GMM objective function for. Specifically, destination fixed effects are estimated with an ordinary least squares (OLS) regression of $\xi(\tilde{\theta}^D)$ on the 119 destination dummies. The residual of this regression, $\tilde{\xi}(\tilde{\theta}^D)$, forms the basis of the GMM estimator of the remaining demand parameters θ^D

$$\hat{\theta}^D = \underset{\theta^D \in \Theta}{\arg\min} \frac{1}{N} g(\theta^D) A g(\theta^D)' \tag{18}$$

$$g(\theta^D) = \tilde{\xi}(\tilde{\theta}^D)' Z^D. \tag{19}$$

$g(\theta^D)$ denotes the empirical moment condition (with destination fixed effects concentrated out), A a positive definite weighting matrix, and N the number of observations. The two-stage GMM estimator of Hansen (1982) is implemented, with in the first stage a weight matrix $A = (\frac{1}{N} Z^{D'} Z^D)^{-1}$ and in the second stage a weight matrix $A = (\frac{1}{N} g(\hat{\theta}^{D1})' g(\hat{\theta}^{D1}))^{-1}$ ($\hat{\theta}^{D1}$ denoting the estimation result from the first stage).

The solution to (18) is obtained with a constrained minimizer using the interior-point algorithm and analytic gradients, with default convergence criteria (1e-6 for the values of the objective function and the constraints). The solver starts from initial parameters that are the estimated coefficients in a nested logit model with one consumer type, obtained from an instrumental variable (IV) regression of $\log(\hat{s}_{jt}) - \log(\hat{s}_{0t})$ on Z^D (see Berry (1994)), and setting starting values of 0.5 for γ_1 and κ. Bounds are set only for $\kappa \in [0.01, 0.99]$, $\gamma \in [0.01, 0.99]$, and $\alpha \in [-10, -0.001]$ parameters, and the results are always in the interior. The standard fixed point algorithm of Berry et al. (1995) is adopted when solving for $\xi(\tilde{\theta}^D)$, with a tight convergence criterium of 1e-12. Overall, the performance of the described implementation of the estimator is surprisingly good given the scale of the problem. Results for the preferred specification are obtained in half an hour on a local machine.[31]

With the estimated $\hat{\theta}^D$ in hand, the marginal cost parameters $\hat{\theta^S}$ are obtained by Instrumental Variable regression. The marginal costs are obtained from (14) as $P_{jt} - B^*_{jt}(\mathcal{F}; \hat{\theta}^D)$ and regressed on W_{jt}, instrumenting for endogenous quantities (in the equilibrium mark-up B^*, see (10)) with supply-side instruments Z^S.

5.2 Identification and Instruments

Identification of θ^D follows from Berry and Jia (2010), and more generally from the results in Berry and Haile (2014) for identification of nonparametric (random coefficient) differentiated product models using market-level data and scalar unobserved heterogeneity. Identification is conditional on having at least one valid instrument for each endogenous variable, satisfying exclusion and completeness restrictions.

The price instruments (Z^D) must generate variation in average flight fares across all flights in a market. The fuel consumption per seat of the flight generates a primary source of such variation. Fuel costs comprise a large share of the costs to operate a flight and vary substantially across aircraft (Brueckner et al. (2023)). I compute a detailed measure using the aircraft-specific fuel consumption parameters of Seymour et al. (2020)

[31] Similar outcomes are obtained when first using the interior-point algorithm with a convergence criteria of 1e-5, and then switching to the sequential quadratic programming algorithm starting from the interior-point results. Both are large-scale algorithms suitable to the size of the problem. Implementing the Merrow-Skerlos algorithm to solve for $\xi(\tilde{\theta}^D)$ does not lead to substantial speed improvements. See Conlon and Gortmaker (2020) for further details about optimization choices and best practices. The performance is substantially worse when estimating demand and supply parameters jointly, which takes at least 15 times more function value iterations depending on the specification.

applied to the flight distance. The fuel consumption in kg per seat is interacted with the Kerosine-Type Jet Fuel price 9 weeks before departure in dollars and adjusted to euros using the exchange rate in that week, to reflect differences in fuel costs over time for flights with the same aircraft-destination-month triple. An identifying assumption is, therefore, that consumers do not value (or observe, at the time of booking) how fuel-efficient the aircraft that they fly with is.

Other aircraft-specific cost shifters in Z^D are the share of seats allocated to first class or business class, whether the plane is a wet lease (owned by another airline), and whether it is large (having more than 150 seats). Other variables affecting the marginal costs of a flight through scheduling efficiencies include the average spacing in hours between flights to any destination by the same airline on the same date and the number of destinations that the airline serves, as well as the averages of these variables across all the carrier's flights departing from the destination, and the number of flights by the carrier departing from the destination. What is commonly referred to as *BLP instruments* are included, too, to capture to some extent the distance in characteristics space of the flight to other flights in the same market.[32] Finally, the number of products in the market and their total capacity (number of seats) is included in Z^D to instrument for the endogenous probability of purchasing *any* flight in the market.

The instruments are remarkably strong. This is especially true for the Fuel Consumption instrument, which already explains 79 percent of the variation in prices. The Wu-Hausman test of endogeneity and a test for weak instruments both support the validity of this instrument (with p-values virtually 0).[33] This remains true when adding the additional instruments in Z^D, when adding destination fixed effects, and when restricting the sample to flights with observed raw fares. In addition, the Sargan test for over-identification rejects that the instruments provide conflicting information about the price coefficient (p-value virtually 0).

The same set of instruments is included in Z^S (although cost shifters are captured by W_{jt} directly), affecting endogenous quantities through the equilibrium mark-up. These mark-up shifters have a strong relationship with quantity. Table C8 in the online appendix lists all instruments and supports the descriptive statistics referenced in this section.

5.3 Model selection

While the outlines of the model have been motivated, which exact model specification best fits the data remains to be determined. This is particularly important to determine for the demand model because consistency and asymptotic normality of the GMM estimator rely on the correct specification of the model and moment conditions (Andrews and Lu (2001)). In addition, how much airlines favor certain slots over others depends on how much consumers prefer to fly at those hours, so the demand model must capture consumers' departure time preferences well.

I rely on the MMSC-BIC criterium of Andrews and Lu (2001) to select the best-fitting specification. The BIC-MMSC is suitable for the selection of non-nested models estimated with GMM. It is based on the Hansen (1982) J test statistic for over-identifying restrictions and penalizes using fewer moment conditions or additional parameters. Let b denote a vector of parameters defined on a bounded support B, c a vector of moments, and

[32] BLP instruments are considered complementary here, as the large market asymptotics may reduce their identifying power if the dependence of the equilibrium mark-up on (rival) product characteristics decreases too quickly (Armstrong (2016)).

[33] I test this in a nested logit model without random coefficients, regressing $\log(S_{jt}) - \log(S_{0t})$ on X_{jt}, P_{jt}, and the product share. The latter two are endogenous variables instrumented with Z^D that includes the number of products per market. See the bottom panel of Table C8 for the results from IV diagnostics tests.

Table 1: Model Selection

	Departure-time preferences specification			
	Hours	Hours (dest. FE)	Cosine Prefs.	2-Type Cosine
MMSC-BIC Model Selection (scaled by N)	0.102	0.037	0.034	0.033
Number parameters	43	157	141	142
Number moments	61	176	160	160
Don't reject H0 correct specification	✓	✓	✓	✓
Computation Time (demand model, hours)	0.141	0.074	0.037	0.088
Stepsize at optimum	0.000	0.000	0.000	0.000
First Order Optimality	0.000	0.017	0.036	0.032
FOC: Max. Abs. Gradient	0.011	0.054	0.121	0.096
SOC: Min. Eigenvalue Hessian	0.000	0.023	0.217	0.219
Share Negative Costs	0.485	0.066	0.043	0.041
Share Inelastic Demand	0.463	0.018	0.007	0.006
Nested Logit Coefficient (kappa)	0.874	0.800	0.780	0.783
Share Business Type	0.623	0.042	0.044	0.045

Notes. The demand models are based on the same model with random coefficients on the constant, the price, and the departure airport. The estimated departure time preferences in the four model specifications are plotted in Figure C5. The checkmarks indicate that for none of the models the null hypothesis that the model is correctly specified rejected at the 5 percent level. The indicator is computed using that $J_n(b,c)$ has an asymptotic chi-square distribution with $k = |c| - min\big(|b|, |c|\big)$ degrees of freedom. The corresponding critical value equals $\chi^2_k(0.05)$, for the 95th quantile of the chi-squared distribution with k degrees of freedom (Andrews and Lu (2001)).

n the sample size. $J_n(b,c)$ denotes the J test statistic

$$J_n(b,c) = n \inf_{b \in B} G_{nc}(b)' W_n(b,c) G_{nc}(b), \tag{20}$$

resulting from the GMM estimator, with $W_n(b,c)$ the corresponding weight matrix and $G_{nc}(b)$ the value of the moments associated with b and c. The MMSC-BIC statistic is defined by

$$J_n(b,c) - \big(|c| - |b|\big) \ln(n). \tag{21}$$

Table 1 reports the computed values of the MMSC-BIC statistic for four different demand specifications, alongside other relevant statistics. The first two columns are based on demand models with departure hour dummies (12), with the benefit of capturing preferences flexibly. The first thing to note is that in the second column, where destination fixed effects are included, the MMSC-BIC statistic is substantially lower.[34]

The estimated departure time coefficients follow an interesting pattern. They peak at 8 am —consumers strongly favoring morning flights, all else equal— and dip about 12 hours later around 8 pm. However, it remains an open question whether travelers' true preferences are as distinguished as represented with the departure hour dummies; do people really like noon *and* 2 pm flights more than flights leaving at 1 pm? I consider it more likely that the true time preferences are more smooth, and that departures from a smooth preference representation result from the model imperfectly capturing other factors that interact with the departure time. This motivates

[34]The estimated departure hour coefficients are plotted in Figure C5 (plot (a) and (b)) in the online appendix. The model with destination fixed effects is preferred because omitting them (plot (a)) removes some of the distinction between the hours, potentially by conflicting preferences for reaching the destination by plane and when these flights are scheduled. In addition, the model without destination fixed effects results in an implausibly large share of negative (net) costs and inelastic demand.

my use of a cosine function as a smooth process for departure time preferences, repeating every 24 hours and forced to peak at 8 am, defined as

$$Cosine24 = cos\big(\frac{\pi}{12}(\text{Departure Hour} - 8)\big). \tag{22}$$

Besides considering a smooth representation (fit) of the estimated departure time dummies, I also consider two models that impose the smooth cosine pattern on the decision-making process. This is done by including the variable Cosine24 as defined in (22) in the utility function in place of the departure hour dummies. The estimated coefficient on this variable is the amplitude of the cosine. The coefficient is interpreted as the utility gain (loss) from the flight departing at 8 am (8 pm) rather than at 2 pm or 2 am — the latter being times when $Cosine24 = 0$. The results of this specification are given in column 3 of Table 1. Another benefit of using the smooth cosine representation is that it can be estimated separately for the two departure type, as in column 4.[35]

Following Andrews and Lu (2001), I systematically consider these different specifications. The null hypothesis that the model is correctly specified is for none of the models rejected at the 5 percent level.[36] As the 2-Type Cosine model has the smallest $|c| - |b|$ and the smallest MMSC-BIC value, it has the best statistical fit of the considered specifications. This specification also has the lowest share of negative (net) costs and inelastic demand and is still estimated fast. The nested logit coefficient (κ) is stable across the three specifications with destination fixed effects. Figure C5 also shows that the departure time preferences of the two traveler types are estimated to be substantially different, and it is a benefit to be able to capture this in the analysis. The resulting model estimates are discussed in the next section.

6 Estimation results

This section presents and interprets the estimated model parameters. The full set of estimated parameters is given in Table C9 in the online appendix.

6.1 Demand parameters

The estimation results indicate that a small share of the population ($\hat{\gamma}_1 = 0.54$) has preferences that diverge from the rest. Based on the estimation results, Type 1 can be considered the business traveller having a lower price elasticity ($\hat{\alpha}_1 = -0.274$ and $\hat{\alpha}_2 = -1.245$) and Type 2 can be labelled as the consumer that travels for leisure. The results confirm that consumers prefer to fly with a large carrier at the departure airport. Still, it does not necessarily come through the number of destinations the airline serves ("Hub"). Instead, consumers value flying with a carrier that offers more flights on the same day to the destination they want to reach. The quadratic term and the scaling of this variable ("Nr. Flights to Destination") imply that a flight is more valuable when the carrier has up to 13 flights to a destination per day.[37] It is plausible that consumers expect more

[35] By contrast, including a random coefficient on the departure hour did not appear feasible.

[36] Specifically, $J_n(b, c)$ has an asymptotic chi-square distribution with $k = |c| - min(|b|, |c|)$ degrees of freedom. The corresponding critical value equals $\chi^2_k(0.05)$, for the 95th quantile of the chi-squared distribution with k degrees of freedom (Andrews and Lu (2001)). The null is not rejected when the (unscaled) MMSC-BIC statistic exceeds the critical value.

[37] The non-linearity in the network effects (coming through the number of destinations or the number of flights to a certain destination) implies that re-allocating a flight from a carrier with a larger airport presence to a newcomer will affect total consumer utility even when holding all else constant. This might reflect the possibility that a carrier with a hub at the departure airport has more to lose when its departure rights are up for auction than that a newcomer has to gain — when abstracting from all other differences between these airlines and their use of the slot.

Table 2: Price elasticities and willingness to pay

Own price elasticities	
Type 1	-1.051
Type 2	-4.791
Both types (population weighted average)	-3.166
Share of passengers	
Type 1	0.045
Willingness to Pay for CDG	
Type 1	-566.170
Type 2	54.398
Willingness to Pay for Connectivity	
Departure Airport - Type 1	222.518
Departure Airport - Type 2	48.958
Arrival Airport - Type 1	77.194
Arrival Airport - Type 2	16.984
Amplitude Cosine	
Type 1	0.081
Type 2	0.204
Smooth departure time preferences	
WTP Departing at 8am vs 2pm - Type 1	29.553
WTP Departing at 8am vs 2pm - Type 2	16.391

Notes. The Willingness to Pay (WTP) is reported in euros per ticket. The WTP for CDG is the Willingness to Pay to depart from Paris Charles-de-Gaulle rather than Paris Orly Airport. The WTP for Connectivity at the Departure (Arrival) Airport captures the monetary value of one additional flight arriving at the departure airport (departing from the arrival airport), operated by the same carrier or a codesharing partner, within 1.5-6 hours before (after) the flight in question is scheduled to depart (arrive). The WTP to depart at 8 am vs. 2 pm is obtained by dividing the estimated type-specific amplitude by (minus) the price elasticity of that type.

delays or cancellations when more than 13 flights are available. Moreover, consumers prefer flights with better connection possibilities at both the departure and arrival airports. Conditional on hub and network benefits, the results suggest that Air France flights are not preferred over flights by Low-Cost Carriers.[38]

The estimated destination fixed effects capture the benefit of flying to that destination compared to reaching it by other means, such as by train. Table C10 reports the 20 destinations with the highest and the lowest estimated fixed effects. As expected, the lowest flight-utility destinations are those that can be reached by (high-speed) train, such as London, Zurich, Milan, and Marseille. The highest flight-utility destinations, such as Atlanta, Saint-Denis, and Dzaoudzi, are all separated from Paris by an ocean.[39]

Table 2 provides further intuition for the estimated demand coefficients, reporting the own price elasticities and willingness to pay for important flight attributes. The price elasticity of the tourist type (-4.79) is comparable to the estimated value in Berry and Jia (2010) (-5.0 in the base specification), and business types are more price elastic (-1.1 vs. -0.44). As only 4.5 percent of passengers are business types in my model, the resulting aggregate price elasticity is twice as large (-3.17 vs. -1.55). Note that the price elasticities in most airline studies (including Berry and Jia (2010)) are based on aggregated data. This likely introduces an upward bias to price-elasticity estimates, as documented in D'Haultfœuille et al. (2022) in the context of rail transportation. Corresponding to this idea, my estimated aggregate price elasticity of -3.17 is close to the -3.31 average price elasticity across markets in Williams (2022), also estimated using flight-level data.

The small share of business travellers are characterised by a strong preference for flying from Paris Orly airport. Their strong preference can be justified because it is more convenient to reach the airport by taxi from

[38] More specifically; the excluded category are non-flag carriers other than easyJet and Vueling Airline.
[39] Saint-Denis and Dzaoudzi are on two French overseas territories in the Indian Ocean, near Madagascar.

central Paris, and it takes less time to go through security and reach the gate. Moving between terminals for flights with a layover at CDG is notoriously complicated.[40] The business type has a high value of time. Also, regarding connectivity, Type 1 is willing to pay €223 for each additional (potentially connecting) flight that arrives at the departure airport within 1.5-6 hours before the flight in question. Also, leisure travelers value flights with better connection opportunities more and are willing to pay €50 for each additional connection at the departure airport. This is important for the counterfactuals; allocating slots from airlines with larger networks (like Air France, who has its global hub in Paris) to airlines who are smaller reduces the consumer convenience of having short layover times.

The estimates also reveal a sizeable departure time preference; business travellers are willing to pay €30 more and leisure travellers €16 more to depart at 8 am rather than at 2 pm. From the airline's perspective, these are substantial departure time premiums, making up about 25-50 percent of the median profit per seat.

6.2 Supply parameters

I estimate that a flight generates on average 22k in profits before fixed costs (see Table 3). Profits are lower for short-haul flights, and distance categories have substantial heterogeneity. The median profits per seat are €64, the average is €110, and at the 90th percentile, each seat generates €309 in profits before fixed costs. Air France is estimated to have similar profit statistics, except for long-haul flights with higher profits (€79k vs €66k on average per flight). On average, the marginal costs make up 53 percent of the price. A Lerner Index of 0.47 on average is lower than the estimate for direct flights in Berry and Jia (2010), consistent with the higher price elasticity uncovered with my flight-level analysis and my rich cost specification, and reflecting relatively tight margins in this industry.[41] The Lerner Index is also 0.47 for Air France flights, matching their self-reports average gross profit margin of 48 percent (across all Air France flights in 2018).[42]

Besides affecting the convenience on the consumer side, the estimation results indicate that the flight network also affects the bottom line through cost synergies (Figure 2 displays the standardized effects of a one standard deviation increase in the variable). I estimate diseconomies of scope where a larger network is more costly for a given number of flights, although there are economies of scale in the number of flights. In addition, carriers face scheduling efficiencies: cost synergies from scheduling flights more closely together. As shown in Figure 2, a one standard deviation increase in the spacing of flights from the departure airport (by the same carrier and date and to any destination) increases costs by 11.7 percent, when keeping the other variables fixed at their sample means. The spacing of flights is the fifth most important cost shifter, after measures of the duration of the flight, the aircraft's fuel efficiency, and the network size in terms of the number of destinations reached. Increasing the number of destinations by one standard deviation results in a 25 percent increase in marginal costs.

Departing from CDG costs about €34 more per ticket, c.p. Differences in fees partly account for this cost difference – the supplemental per-passenger fee being, for instance, more than twice as high at CDG than at

[40]On the other hand, high-end lounges and airport facilities are inferior at ORY. Anecdotally, according to the Points Guy website, *Charles de Gaulle can be an infuriating airport to navigate depending on what terminal you're departing from, but once you clear security you'll have many more options for your time, be it shopping, eating or relaxing in a lounge.*

[41]Other Lerner Index estimates in the literature range between 0.32 and 0.39, with lower estimates for direct flights in a model without random coefficients (Bontemps et al. (2021)) and higher values for integrated codeshare products (Gayle (2013)), all for US markets in recent years and estimated based on disaggregated DB1B data.

[42]Source: Macrotrends.net, who provide profit statistics based on financial reports of traded companies. Air France's financial statement over this period provides further cost breakdowns.

Figure 2: Marginal cost parameters: standardized effects

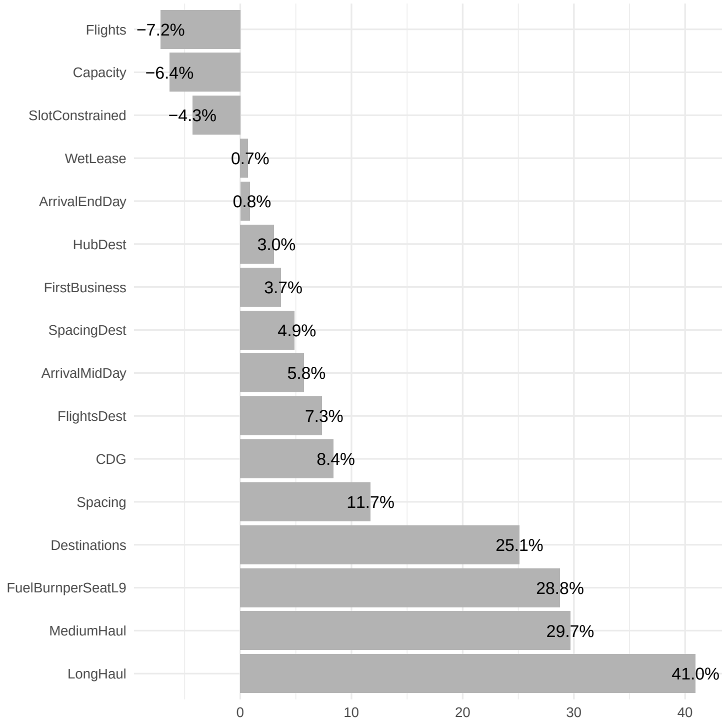

Notes. Displaying the percentage change in marginal costs when increasing the relevant variable by one standard deviation, holding the other variables at their means. For binary variables, the effect measures the percentage change in costs when setting them from zero to one. All estimated coefficients are statistically significant. See Table C9 for the corresponding point estimates.

ORY and connecting passengers at CDG being even more expensive.[43] As expected, the average cost per ticket on flights with a higher business class share is more expensive, as well as those with a higher fuel consumption per seat, on a plane leased from another airline, and for flights that are medium or long haul. Arriving in the middle of the day is estimated to be costly, which might reflect the network benefit of having the aircraft operational for a return flight at a reasonable hour. As mentioned above, the marginal cost should be interpreted as the average variable cost per seat net of any network benefits. For 4.1 percent of flights, I estimate negative costs, which might reflect feeder flights for which the benefit of routing passengers through the destination outweighs the flight operation costs.[44] Relatedly, the cost per seat for flights to slot-constrained destinations is about €17 less, which can similarly be interpreted as those flights incurring a large benefit of having the aircraft arrive at the destination.

7 Estimated slot values

Section 7.1 discusses the slot values derived in the data. The role of the airline's flight network and the related magnitude of departure slot complementarities is further investigated in Section 7.2.

[43]Source: Fee schedules ORY and CDG for 2018.

[44]Empirically, the probability that costs are negative is higher for flights towards a Hub (serving more destinations by the same carrier) in a Probit specification that controls for destination fixed effects.

Table 3: Profits by flight and seat

	All airlines	Air France
Lerner Index		
Average	0.471	0.468
Median	0.404	0.378
Profit per flight (10k)		
Average	2.246	2.300
Median	1.307	1.295
Profits per flight (10k), median by departure airport		
CDG	1.258	1.225
ORY	1.365	1.377
Profits per flight (10k), average by distance		
Short Haul	1.206	1.192
Medium Haul	4.564	3.929
Long Haul	6.626	7.868
Profits per flight (10k), 90th percentile by distance		
Short Haul	1.635	1.667
Medium Haul	8.211	8.635
Long Haul	12.593	14.095
Profits per seat		
Average	110.434	112.208
Median	63.394	65.319
90th percentile	308.581	308.581

Notes. 63,793 out of 120,039 flights are operated by a carrier in the Air France group. Profit statistics for these flights are given in the column labeled *Air France*.

7.1 Slot value results

To assess the welfare impacts of introducing a market for slots, slot values are only derived for the four largest firms; Air France, Lufthansa, easyJet, Vueling Airlines. These four airlines combined operate 60% of flights and sell 68% of tickets in the data. All other carriers make up a small market share and excluding them from the slot auction —not from the market— reduces the computational burden. Specifically, how much the big four airlines (a) value the 252 unique hourly departure slots (s) is determined following Assumptions 1-2. First, the flight whose departure rights expire (f^X) is obtained by selecting a flight randomly among all flights departing in s and picking the flight with the lowest baseline profits among all flights of the same carrier departing in s. Then, each airline computes the incremental profits *across their flight network* when adding the flight with the highest baseline profit to their network and letting it depart at s, subject to the constraints given above.[45]. This requires updating the exogenous variables (X, W) to account for network effects and simulating the market equilibrium in all markets where are least one of these variables changes for at least one of the carriers. The change in total profits gives $\nu(a, s)$.

Computationally, the market equilibrium does not always converge within a reasonable time for all 119 markets, 252 slots, and all carriers operating from the departure airport associated with the slot. Simulations are halted when the price does not converge to within €5 within 500 iterations in at least one of the markets. Slot values are excluded for all bidders if for at least one of the bidders the simulations are halted prematurely. The results are based on the 216 slots (auctions) that converged quickly according to these criteria. To reduce the dependence on outliers, slot values exceeding the 95th percentile are excluded from the results. Furthermore, $\nu(a, s)$ is derived for a random ordering of s to ensure that the constraints' slackness does not relate

[45]Each flight can be added at most once and the overnight status of the added flight has to match the overnight status of f^X.

Table 4: Estimated slot values (in €100k, S18 flight season)

	N	Median	SD	Total
Slot Values	595	5.790	5.874	4708.631
Slot Values - by Bidder				
Air France	166	11.271	8.178	2227.362
Lufthansa	76	6.312	2.917	528.263
Vueling Airlines	143	4.792	2.304	771.987
easyJet	210	5.156	2.150	1181.019
Slot Values - by Airport				
CDG	381	5.949	6.859	3332.949
ORY	214	5.646	2.975	1375.683
Slot Values - Winning bidders	205	12.056	12.208	3386.921

Notes. Slot values across 216 auctions for which all simulations converged, excluding 9 auctions with negative values and with values exceeding the 95th percentile.

systematically to the departure time of the slot.

The 595 resulting slot values are summarised in Table 4. They are the basis of the auction results in Section 8. One departure slot used for a weekly flight in one flight season is valued at €579k, at the median. While the value of a slot is about the same in CDG as in ORY, there is a larger variation across departure hours at CDG airport. This is because some CDG departure slots are used for highly profitable long-haul flights with wide-body aircraft. The dominant carrier has substantially higher (median) slot values than the other carriers. Unsurprisingly, Air France values some slots at its global hub more than the other carriers. Differences in slot values are partly driven by differences in costs and demand, with the denser network and the higher flight frequency (to a destination) generating an advantage for Air France. In addition, Air France operates some of the most profitable flights from Paris to overseas destinations and, when expanding, is expected to offer more of such flights.

When holding the network variables fixed, there is no particular reason why Air France flights generate higher profits or why this carrier should have higher slot values. Quite the opposite, in fact, as a negative Air France premium is coming through the demand side. It is, therefore, a general result that the dominant carrier will have a favorable position in a potential slot auction in its hubs, and, for instance, it is to be expected that Lufthansa values slots at Frankfurt and Munich more than Air France does. Existing network asymmetries will, to some extent, carry through in the allocation of a market-based mechanism unless it is designed to favour small bidders or newcomers particularly.

Figure 3 externally validates the estimated slot values. Despite the unavailability of systematic data about prices paid for (or bids placed on) slots, the information available about slot prices at Heathrow Airport aligns with my results. Representing different airports and periods, and in the case of the sample of paid prices also being for a non-random selection of slots, a perfect match is not to be expected. However, one would expect slot values to be somewhat comparable as London Heathrow, Paris Charles de Gaulle, and Paris Orly are all among the most slot-constrained airports in the world, they are close in proximity, and they have similar competitive environments with dominant carriers owning a large share of slots. The figure shows that the model's price estimates for Paris have the same order or magnitude as (estimated) prices for London Heathrow. This is supportive of my approach—recovering potential bidder values based on a model of the market's demand and supply equilibrium rather than a model of equilibrium bidding—which yields reasonable results.

Next, I assess what makes some airport slots more valuable than others. Slot value determinants are

Figure 3: Validation based on price estimates for Heathrow slots

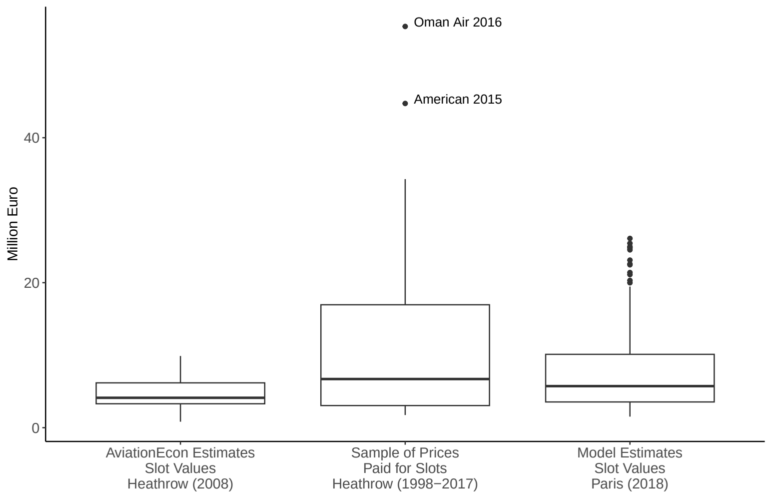

Notes. The Heathrow slot value estimates in the box-plot on the left are from a 2008 newsletter publication by the "strategic aviation consultancy" firm *Aviation Economics*, available upon request. Both their low value and high value estimate for each departure hour are included in the plot. The sources for Heathrow slot prices in the middle box-plot are, for 1998-2013 this CAPA article, and for 2013-2017 Odoni (2020, Table 3.4). All prices are given in 2018 euros for a daily slot pair for 31 weeks (as in S18). The model estimate of a slot price in Paris is the present value of an ordinary annuity that pays $\nu(a, s)$ per year for 5 years (as in the mock slot auctions described in Ball et al. (2007)) with an interest rate of 1.17% per year – the average 2018 inflation rate in the European Union (Macrotrends)), taking only $\nu(a, s)$ for the winning bidder of each slot (the highest across a).

Table 5: Slot value determinants

	coefficient	std. error	marginal effect
Number competing flights in market	0.488*	(0.205)	11.33
Herfendahl Index	1.986**	(0.715)	6.295
Distance of flight (1000 km)	1.763***	(0.273)	54.747
Capacity of aircraft (100 seats)	3.645***	(0.294)	57.774
Connectivity arrival airport	0.953*	(0.392)	7.578
Connectivity departure airport	-2.557***	(0.692)	-29.137
Morning slot (before 11 am)	1.364***	(0.409)	25.785
Overnight flight	-3.345*	(1.386)	-63.229
Dominant carrier in global hub	2.677***	(0.687)	50.599
Slot owner	1.413***	(0.423)	26.719
Slot at CDG	-0.464	(0.6)	-8.781
Weekday slot	-1.234*	(0.556)	-23.335
Weekday slot at CDG	2.026**	(0.72)	38.291
Intercept	-4.508***	(0.921)	

Notes. The coefficients and standard errors are from an OLS regression with the slot value in €100k as a dependent variable in the sample of 595 estimated slot values summarized in Table 4. The stars indicate the following statistical significance levels: * 0.1, ** 0.05, and *** 0.001. The last column reports the standardized effect. This is computed as the percentage change in the slot value when the numerical (binary) variable changes by one standard deviation (from zero to one), holding the other numerical (binary) variables at their sample means (at 0).

obtained by regressing the estimated $\nu(a, s)$ on variables relating to the slot, the carrier's network, the flight, and competition at the route level. The estimated coefficients and standardized effects are given in Table 5. As expected, morning slots are more valuable. Departing before 11 am increases the slot value by €14k (25 percent). Slots at CDG are more valuable than those at ORY, but only during weekdays. Using the slot for a flight in a more concentrated market also increases the value. All else equal, slots are more valuable to carriers that currently own them than to carriers that compete for them, suggesting the benefit of optimizing their flight

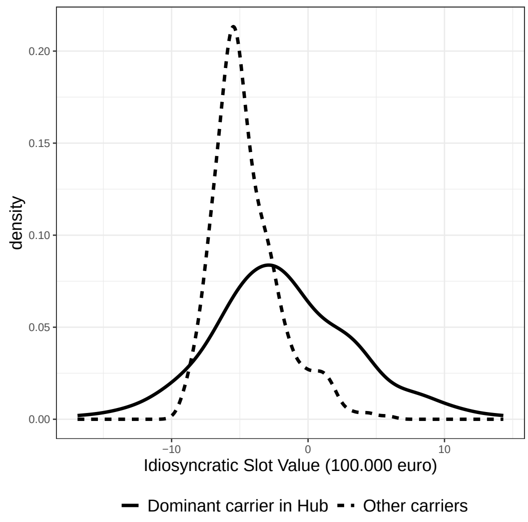

Figure 4: PDF of idiosyncratic slot values, $f_{V|\mathbf{A}}$ with X the determinants in Table 5. Specifically, the idiosyncratic slot value is computed as the residual from the regression of $\nu(a, s)$ on $\mathbf{X}$, adding the intercept and, where relevant, the global hub effect.

network around the slot. In addition, also reflecting network benefits, slots are more valuable to Air France than to the four other carriers. Interpreting this in general terms, slots are estimated to carry a large premium for the dominant carrier in their global hub, who values a slot €267k more than the other firms. The dominant carrier benefit is of the same magnitude as using the slot for a flight with one standard deviation larger capacity or for a destination one standard deviation further away. Additional slot value determinants are given in the table.

One also requires the distribution of *idiosyncratic* values to simulate bidding in a specific auction. These are the $\tilde{V}_i$ in (7). Assuming additive separability so that $V_a = h(\mathbf{X}) + \tilde{V}_a$, the idiosyncratic values are obtained as the residual from the regression of $\nu(a, s)$ on $\mathbf{X}$, with $\mathbf{X}$ the variables reported in Table 5, adding the intercept and for slot values of Air France adding the coefficient on *Dominant carrier in global hub*. The resulting density of idiosyncratic values, $f_{V|\mathbf{A}}$, is plotted in Figure 4. Consistent with the analysis above, an auction model with asymmetric bidders is best suited for this market. The first two moments of $f_{V|\mathbf{A}}$ are higher for a carrier with a large airport presence in one of its hubs than for the other carriers.

7.2 Complementarities and network effects

I also use the model estimates to quantify the magnitude and determinants of departure slot synergies.[46] To further illustrate how flight networks affect slot values, I focus on three market settings; a small carrier in a small market (with a low flight frequency), a large carrier in a small market, and a large carrier in a large market. These settings are represented in the data by Aeromexico (the small carrier) operating one daily flight to Mexico City in September (the small market), Air France (the large carrier) operating one daily flight to Mexico City in September, and Air France operating 13 daily flights to Toulouse in July (the large market), respectively.

The simulations subsequently have the carrier win 1, 2, 5, and 10 additional slots for these three settings.

[46] As noted earlier, I consider daily slot pairs as the unit of analysis, which already capture obvious complementarities between arrival and departure slots in one airport and which are sufficient even for existing carriers to expand to a new destination (see Table C2). This is also the relevant unit when slot trades occur and the one that would likewise reduce the exposure problem in a combinatorial auction in one airport alone.

Table 6: Departure slot packages
Values in €100k, S18 flight season, and as share of value for first slot

			Aeromexico-MEX		AirFrance-MEX		AirFrance-TLS	
	updating	slots	value	share 1st	value	share 1st	value	share 1st
Only updating network variables	All 3	1	4.852	1.000	0.962	1.000	0.200	1.000
	All 3	2	14.100	2.906	1.732	1.802	0.220	1.096
	All 3	3	18.933	3.902	2.434	2.531	0.188	0.936
	All 3	4	21.999	4.534	3.055	3.177	0.103	0.514
Updating + adding flight to market	All 3	1	20.200	1.000	25.320	1.000	3.513	1.000
	All 3	2	65.346	3.235	51.266	2.025	6.980	1.987
	All 3	3	109.136	5.403	77.507	3.061	10.356	2.948
	All 3	4	150.320	7.442	103.708	4.096	13.596	3.871
Contribution of Individual Network Variables								
Only updating network variables	Spacing	1	4.526	1.000	0.353	1.000	-0.002	1.000
	Spacing	2	13.096	2.893	0.617	1.749	-0.094	46.884
	Spacing	3	17.293	3.821	0.919	2.607	-0.148	74.218
	Spacing	4	19.714	4.356	1.256	3.561	-0.172	85.779
Only updating network variables	Flights	1	0.000	1.000	0.353	1.000	-0.003	1.000
	Flights	2	0.000	1.143	0.617	1.749	-0.097	34.252
	Flights	3	0.000	1.000	0.919	2.607	-0.156	55.008
	Flights	4	0.000	1.143	1.256	3.561	-0.185	65.231
Only updating network variables	Frequency	1	0.231	1.000	0.962	1.000	0.201	1.000
	Frequency	2	0.423	1.833	1.732	1.802	0.223	1.108
	Frequency	3	0.574	2.487	2.434	2.531	0.196	0.971
	Frequency	4	0.682	2.955	3.055	3.177	0.117	0.582

Notes. Slots are complements when the value of adding N slots is larger than N times the value of the first slot. The presence of complementarities thus requires that *share 1st* > *slots* in the rows that are based on the full equilibrium where the flights are added to the market and where all 3 network variables are updated.

After updating the relevant network variables in (X, W), the new equilibrium is computed across all markets, and the value of the slot pair is taken to be the difference in the carrier's total profits. To set aside any (selection) effects associated with the choice of flight, these simulations restrict both the number and use of the slots won by bidders; all flights are out of CDG on the first Monday of the month at 9 am. The flight characteristics at baseline (X, W, ξ, ω) are from the (nearest) flight already operated on that day by the carrier.

The results are given in Table 6 and paint a nuanced picture. There are two ways in which airlines can benefit from slot package complementarities. First, package slots can be more valuable when offering more flights to the same destination, as consumers value frequency. This holds up to a point; the peak is reached at 13 daily flights given the estimated model parameters. As this is exactly the baseline number of flights in the large carrier / large market example, with Air France operating 13 flights per day from CDG to Toulouse in July, this contributes to the pure network effect of adding the 14th flight being negative (amounting to -6k, see the second-to-last column in the first row of Table 6). The second way package slots can be more valuable is when they reduce the operational costs of flights departing on the same date, regardless of the destination. The contribution of the three network variables is separated in the lower part of the table.

The benefit to Air France of adding the first flight to Mexico City (the large carrier / small market example) is substantial (25k). Slot values like this are at the upper end of the slot value distribution because they hit a sweet spot: combining the benefits of the large carrier, adding a profitable long-haul flight to a low-frequency market, and banking on cost synergies across many flights in its network. The results indicate that higher demand due to the increased flight frequency within the market drives most of the network benefits. Scheduling efficiencies through the number of flights or their scheduling for departures on the same date make up less than 25 percent of the network benefits in this setting. In addition, the contribution of the network benefits is relatively small, making up only 6 percent of the total value of the first slot. To conclude that network effects are minor would be misleading, however; Air France values the first extra slot to Mexico City 2.5 times more than Aeromexico mostly because of the size of Air France at the departure airport.

Another striking result is the magnitude of the network effects in the small carrier / small market example. When adding two additional flights to Mexico City, all network effects combined generate twice as much profit for Aeromexico as for Air France. In fact, this leads to the result that the small carrier / small market case is the only one with slot package complementarities. In the other two cases, any network effects are fully undone in the full market equilibrium when the flight is added to the network, and no benefit of winning multiple slots at once is detected. Generally, departure slot packages allow for the enjoyment of the benefits of a larger network (coming from demand premie and cost savings), which are valued more by smaller airlines for whom the expansion is larger in relative terms. We can conclude from the simulation exercise that departure slot packages can benefit smaller carriers more than larger ones.

8 Introducing slot auctions

Before quantifying the welfare impacts of introducing a market for slots in Section 8.3, Section 8.1 outlines considerations about the form of such a market.

8.1 Auction design

An efficient market is implemented in what follows by allocating each s to the carrier with the highest $\nu(a, s)$. This is done according to the model's assumptions and based on counterfactual market simulations that account for spillovers across flights and markets, as detailed in Section 7.1. The focus lies on the implications of a slot market on the competitive environment, social costs, and consumer welfare rather than on the allocation of auction rents between bidders and the auctioneer. Any transaction prices resulting from the auction do not affect the total welfare as it is credited to the auctioneer (the airport or the slot coordinator), and, being a sunk cost at the flight level, do not affect equilibrium fares. One may consider —for political reasons, perhaps— the redistribution of auction revenues to the market participants by reducing landing rights and other airport fees.

The auction is for series of Weekly Slot Pairs (WSP), consisting of a departure and matching arrival slot in the origin airport for a certain time each week in the flight season. The value of all slots is computed independently, considering the added value of the slot to the carrier's baseline network of flights. Section 7.2 highlights that synergies of adding multiple WSP are small at best, except for the marginal carriers that can benefit from unlocking cost synergies across multiple additional flights. In those instances, a counterfactual expansion of the carrier's network is also the most speculative, requiring aircraft availability and demand for return flights. This motivates me to consider only the biggest four carriers competing for WSP at auction (another motivation is computation time). The setting is one with independent private values.[47] It is assumed that there are no borrowing frictions so that bidders do not face budget constraints. Many auction designs

[47]Note that if the value definition (6) would include the *probability* that competitor a' wins the auction, which would be a function of the value that a' attaches to the slot, bidder values would be interdependent by construction.

Table 7: Auction results

	Pre-auction	Value definition according to model: Base case	Add quantity cap	More participation	Alternative value def.: No hoarding	Drop random flight
Expected revenue (per slot, 1m)		2.432	2.238	27.347	2.280	2.349
Winner surplus (per slot, 1m)		2.930	0.804	96.223	2.414	2.586
Number bidders		3.000	3.000	13.000	3.000	3.000
Share slots owned by 4 firms	0.648	1.000	1.000	0.574	1.000	1.000
Proportion owned by						
Air France	0.736	0.741	0.362	0.574	0.679	0.647
Lufthansa	0.050	0.051	0.157	0.000	0.091	0.034
Vueling Airlines	0.050	0.056	0.200	0.000	0.105	0.092
easyJet	0.164	0.153	0.281	0.000	0.124	0.227

Notes. The reported expected auction revenue is the median of the second-highest value ($\mathbb{E}[V_{N-1:N}]$). The winner surplus is the median difference between the highest and second-highest value ($\mathbb{E}[V_{N:N} - V_{N-1:N}]$). The number of bidders (N) is the median number of carriers for which the bidding constraints ($\mathbb{C}_s$ in Assumption 1) are not binding in the auction of slot s, meaning that these carriers have operations out of the origin airport with the same overnight status as the flight whose rights expire (f^X) and that these flights have not yet been added to the carrier's flight network in previous slot auctions. Further details are given in Section 8.2.

efficiently allocate independent items to IPV bidders, including first- and second- price sealed bid, English, and Dutch auctions. By the revenue-equivalence theorem, these auctions yield the same expected revenues for the auctioneer, and in the second-price versions, it is a dominant strategy to bid truthfully (Krishna (2009)). In what follows, one can consider any of these auction designs as a basic auction with the above-mentioned properties (referred to as *base case* in the results).

Realistically, the actual auction design would need to be tailored to the specifics of the market. The Clock-Proxy Auction (or CPA, developed by Ausubel et al. (2006)) is probably the most suitable practical auction design for this context and has been proposed for auctioning slots at NYC La Guardia Airport based on good mock auction results (see Ball et al. (2007) for details).[48]

8.2 Auction results and alternative slot values

Table 7 presents the first set of results. Before the auction, the four auction participants own two-thirds of the auctioned slots. The slots are allocated to the four bidders as follows: Air France has 74 percent of the departure rights, easyJet has 16 percent, and Lufthansa and Vueling Airlines each have 5 percent. One interesting finding is that the auction participants grow proportionally while owning more slots after the auction. In particular, the column labeled *Base case* in Table 7 shows that when allocating the slots through an efficient mechanism, Air France retains 74 percent of the slots in the auction, and the others also keep their relative sizes except for a slight redistribution of slots from easyJet to the other three bidders.

The resulting slot ownership proportions are similar, but many slot trades take place, resulting in a large surplus for the winning bidder of on average €1.1 million per slot. This is computed as the average difference between the highest and second-highest value ($\mathbb{E}[V_{N:N} - V_{N-1:N}]$). The absence of a secondary market for slots reconciles these results; carriers cannot generally take advantage of opportunities to sell their least valuable slots

[48] Even if bidders struggle to determine their true values, Parkes and Ungar (2000) demonstrate that iterative combinatorial auctions lead to allocative efficiency and offer significant computational advantages over their closed-form equivalents. Finally, it is worth noting that the proposed CPA design achieves a "core allocation" relative to reported preferences (Ausubel et al. (2006)), with desired properties of efficiency and revenue maximization (Day and Milgrom (2010)).

and purchase slots that would add more value to their network. It is generally believed that slots are not traded within the EU because of unclear regulation and ill-defined property rights.[49] The effect of the two policy rules (the quantity cap and the set-aside that enhances participation) are discussed in Section 8.4. The auction results in €2.4 million expected revenues per slot, computed as the average of the second-highest value ($\mathbb{E}[V_{N-1:N}]$) across the 216 simulated slot auctions. A key difference between a well-functioning secondary slot market and an efficient primary market is that in the latter case, the auctioneer receives these proceeds, avoiding the issue of windfall profits in a secondary market.[50] The proceeds can be redistributed to firms and/or consumers by, for instance, improving the airport infrastructure, reducing airport fees, or reducing taxes. The magnitude of the auction proceeds depends on the auction frequency, the duration of the departure rights, and the share of slots that are auctioned off. I assume that the regulator would auction off only 10 percent of the slots each time, to maintain continuity, and that the winning bidder would be given the the slot departure rights for 5 years. In that case, each auction is for about 900 WSP slot series for the two airports in the study. This is estimated to generate €2.2 billion in auction revenues. For context, the auction proceeds would be of the same magnitude as those from the latest auction for 310 MHz of spectrum in the 3.4 - 3.8 GHz band, conducted by the French communications regulator Arcep.[51].

The outcomes of the base case where only the largest four firms participate resemble actual expansion patterns in the market. Information about slots requested by carriers provides a useful source of validation. The 2019 bankruptcy of the airline carrier *Aigle Azul* made an unusually large number of slots available for reallocation at ORY.[52] Of the six carriers that are both in my data and that requested slots from this pool, the two smallest ones that are not projected to win slots at auction also did not receive any slots from the French slot coordinator COHOR (this regards Air Caraibes and Ukraine International). Three carriers that revealed a desire to grow and received COHOR-allocated slots are also projected to win at auction (easyJet, Air France, and Lufthansa). The fourth carrier projected to win at auction indicated a willingness to expand substantially (Vueling Airlines requested 51 daily slot pairs), although it did not get any assigned slots. The projected growth of Vueling Airlines is consistent with receiving all 18 daily slot pairs at ORY, which Air France was forced to give up in 2021 to remedy competition distortions after receiving financial support from the French state.[53] These results speak to the model's ability to characterize market-based outcomes based on an approximation of the relatively complex flight choice model with spillovers across the network.

Table 7 also reports the auction outcomes based on two alternative definitions of $\nu(a, s)$. The first alternative shuts off the effect of slot hoarding. Specifically, in the no-hoarding slot value definition, carriers only consider their expected profits when winning the slot, as in

$$\Big(\sum \mathbf{K}^*_a(\mathcal{F}^{s(a)}) - 0\Big), \tag{23}$$

[49] Odoni (2020) reviews this issue and concludes: *Secondary trading, as practiced today at some Level 3 airports in the UK, has important positive and negative impacts. The realistic options for the future are at this point (i) continuation of the existing hands-off policy that leaves it up to Member States to decide whether to permit secondary trading at Level 3 airports in their territory (to date, only the UK does), or (ii) amending Regulation 95/93 to permit secondary trading throughout the EU, subject to certain conditions.*

[50] Moreover, as argued by Krishna (2009), a regulator should aim for the most efficient primary allocation mechanism, as issues like transaction costs or thin markets can reduce the efficiency of resale.

[51] These Ascending Clock Auctions raised €2.9 billion according to this press release, which would be the equivalent of the approximate revenues from auctioning 13 percent of the roughly 9,000 available WSPs.

[52] For details, see https://www.cohor.org/2019/12/05/ory-05122019-new-slot-pool-allocation/.

[53] The assignment was based on a beauty contest, prioritizing proposals by carriers already operating at ORY and ranking them based on the airline's capacity and the connectivity potential of the carriers' use of slots. For details, see https://ec.europa.eu/commission/presscorner/detail/en/ip_21_4805.

Table 8: Validation based on 2019 Aigle Azul bankruptcy

Aigle Azul bankruptcy slot pool			
Slot-requesting carriers also in data	Slots requested	Slots awarded	Projected to win auctions
AirCaraibes	20	0	FALSE
easyJet	21	7	TRUE
AirFrance	60	14	TRUE
UkraineInternational	7	0	FALSE
VuelingAirlines	51	0	TRUE
Lufthansa	14	14	TRUE

Notes. Based on slot requests and awards based on data from COHOR (available here). *HOP!* is coded as *AirFrance*, representing the Air France Group in this paper. The table lists only the six carriers that vied for Aigle Azul bankruptcy slots and are also in the dataset used for the empirical analysis.

Table 9: Welfare effects of slot auction across markets (in €100k)

	obs	sum	mean	sd
Welfare Change				
Consumers + firms + auctioneer	879	961.163	0.791	21.397
Compensating Variation				
All markets	879	-696.137	-0.792	20.880
Variable Profits				
All bidders	3516	296.847	0.084	0.743
Variable Profits by firm				
Air France	879	225.079	0.256	1.236
Lufthansa	879	17.305	0.020	0.430
Vueling Airlines	879	39.327	0.045	0.594
easyJet	879	15.135	0.017	0.320
Non-bidding firms	94	-31.821	-0.339	0.493

while still having the best alternative flight $f_a(s)$ and the dropped flight $f^X(s)$ determined as in assumption 1. The results underscore that strategic slot hoarding is important. Owing to lower implied values, Air France would win 6 percentage points fewer slots if it were not accounting for the allocative externality of its competitors using the slots in the downstream market. In other words, the results replicate that slot owners are incentivized to *sit on slots* even if they are not profitable individually. The 80 percent use-it-or-lose-it rule contributes to the inefficient use of available airport capacity previously documented.

The second alternative shuts off the effect of slot owners removing the lowest-profit flights from their networks when needing to give up a slot. Hence, the slot value is still given as in assumption 2, but $f^X(s)$ is chosen at random from all flights taking off in the departure slot window, as in

$$f^X(s) = j \in_{\mathbf{R}} \{j \in \mathcal{F}_a | \mathbb{X}_j \in \mathbb{D}_s\}, \tag{24}$$

where $\in_{\mathbf{R}} \{\dots\}$ denotes selection at random among all elements in the set. The results highlight that the choice of which flight to cancel is relevant. When shutting off the ability of carriers to optimize within their networks, the flag carriers Air France and Lufthansa win fewer slots.

Table 10: Welfare effects of slot auction — with quantity cap (in €100k)

	obs	sum	mean	sd
Welfare Change				
Consumers + firms + auctioneer	889	514.346	0.300	10.757
Compensating Variation				
All markets	889	-318.958	-0.359	9.923
Variable Profits				
All bidders	3556	219.838	0.062	0.759
Variable Profits by firm				
Air France	889	95.943	0.108	1.065
Lufthansa	889	3.970	0.004	0.118
Vueling Airlines	889	87.050	0.098	0.947
easyJet	889	32.876	0.037	0.505
Non-bidding firms	80	-24.450	-0.306	0.461

8.3 Welfare effects

In this setting, the efficiency of the market mechanism does not necessarily result in each slot reallocation being beneficial for consumers.[54] I address this issue for the slot auctions in my data. Table 9 reports the overall welfare effects of introducing the auction, taking into account its effect on the resulting equilibrium in the *after-market*, and considering variation in production technologies (W, ω) and sunk assets (represented by $\mathcal{J}$) across competing firms. The results indicate that auctioning the 216 slots increases the total welfare by roughly €96 million. This is computed by assigning all slots to the highest bidders, updating the network variables in (X, W), and recomputing the equilibrium once again across all markets. The results for the final allocation include markets for which the change in simulated profits is within the 1st-99th percentile for all firms, as to minimize the effects of outliers. The reported compensating variation gives the monetary compensation consumers need to receive to be indifferent between the pre- and post-auction scenarios. This measure reflects changes in both prices and products (as in, e.g., Wollmann (2018)).

Overall, consumer surplus increases by €70 million. Due to the modeled spillovers through network effects coming through demand and supply factors, consumers in all markets are affected regardless of whether the product offering changes in those markets. The large standard deviation across markets also points to the distributional impact of the policy. That the variable profits increase for all firms suggests, in part, that the slots are better utilized. The universal increase in variable profit is also partly driven by the fact that only the largest four carriers participate in this auction variant, so they are taking away value-creating slots from non-participating bidders. The profits of non-participating bidders decrease by €3 million in total, which is less than €100k per carrier. These losses are small because slot owners choose to cancel the least profitable flight among all flights in the slot that is targeted to expire.

[54] The key issue with auctioning product licenses rather than products or services is clearly set out in Borenstein (1988): *The potential for inefficient allocation arises because different firms would, if assigned a license, choose to enter different markets or to use different production technologies. A firm's most profitable use of a license depends on its production set and its sunk assets. Sunk assets might include reputation or acquired expertise in a specific industry or market, as well as the machinery and equipment that are costly to trade. Of course, if these factors were identical for all potential license holders, then any applicant would choose the same market and production process. If that were the case, assignment of operating licenses would not affect resource allocation and would probably not be a* ***significant public policy issue*** [emphasis added]. *Such an assumption, however, is equivalent to supposing that a certain landing slot at La Guardia Airport would be used to serve the same route whether the slot were assigned to United Airlines or Air Vermont. In reality, the most profitable use of a license will differ among firms, depending in large part on the products that they already sell.*

8.4 Quantity caps and set-asides

A critical concern in auctions with aftermarkets is the potential for winners to amass an excessive share of items, leading to market monopolization. To mitigate this, a cap on the maximum share of slots each bidder can secure may be imposed, although this could compromise efficiency and inflate costs. I evaluate the welfare implications of such a policy by capping each firm's slot winnings at 30 percent. The 252 slots are auctioned in a random order. Bidders are excluded after attaining the highest $\nu(a, s)$ in 75 slot auctions to enforce the quantity cap by setting their values to zero. The resulting slot allocation, presented in the column labeled *Add quantity cap* in Table 7, may show proportions exceeding 0.3 per bidder due to excluding outlier values and auctions where the market equilibrium did not converge for at least one of the firms in at least one market.

As anticipated, the quantity cap is binding only for Air France. Slots previously held by Air France are now proportionally distributed among the remaining three bidders, resulting in a twofold increase in variable profits for competing low-cost carriers, Vueling Airlines and easyJet, compared to the auction without the cap (see Table 10). Lufthansa is projected to prefer the outcomes of an uncapped auction, despite winning fewer slots, owing to allocative externalities. Implementing the quantity cap halves the monetary gain for consumers, with compensating variation decreasing from €70 million to €32 million. This reflects consumers' high valuation of the convenience associated with the dominant carrier's denser flight network, network-related efficiencies curbing price increases, and disparities in product choice. Consistent with a long literature underscoring the benefits of the hub-and-spoke model in the airline industry, these results highlight the unusual implications for an optimal slot auction design.

Set-asides and subsidies are additional policy tools often considered to achieve distributional goals (Athey et al. (2013)), including the mitigation of downstream market power (Cramton et al. (2011)). Athey et al. (2013) demonstrated that incentivizing smaller bidders through subsidies is preferable to set-asides, as the latter restricts competition. To gauge the impact of increased bidder participation, I simulate an auction where all carriers participate rather than just the top four. In this scenario, Air France secures 57 percent of slots, less than the 74 percent in the base case but more than the 36 percent with the quantity cap. However, this is Air France's least favored policy, which must be due to smaller carriers expanding into more profitable markets for Air France. Similarly, consumers prefer this scenario the least. In terms of magnitude, implementing a quantity cap reduces the gain in consumer surplus by 55 percent. In comparison, a scenario where 43 percent of slots are won by small carriers (reflecting a significant bidder-subsidy program) reduces the gain by 95 percent. Despite these reductions, the overall consumer surplus and total welfare gains remain positive.

When promoting increased participation from carriers, Air France loses slots to some of them, like Air Caraibes, Royal Air Maroc, and Tunis Air.[55] Some new participants are simulated to grow substantially; Iberia for instance had only two WSPs due to expire and leaves the auction with 23 WSPs. Conversely, easyJet loses all the 26 WSPs it entered the auction with.[56] It is interesting to contrast this finding with the fact that easyJet is estimated to gain from a quantity cap, meaning that, to some extent, larger non-dominant firms can benefit from the curbing of market power of the dominant firm, but encouraging too much participation might hurt them, too.

[55] Cross-tabulations of slot ownership before and after auctions are provided in Tables C12 and C13 for the basic and expanded bidding scenarios, respectively.

[56] It should be noted that the set-aside policy simulations extend the most beyond the status quo, and unmodeled factors such as fleet restrictions or budget constraints are likely to curb the projected growth of smaller carriers in this scenario.

Table 11: Welfare effects of slot auction — with small firms participating (in €100k)

	obs	sum	mean	sd
Welfare Change				
Consumers + firms + auctioneer	382	20.263	0.053	0.076
Compensating Variation				
All markets	382	-15.965	-0.042	0.066
Variable Profits				
All bidders	21774	4.298	0.000	0.003
Variable Profits by firm				
AirFrance	1146	5.246	0.005	0.010
FlagOther	14898	-0.527	0.000	0.001
Lufthansa	382	-0.019	0.000	0.000
NonFlagOther	4584	-0.149	0.000	0.001
VuelingAirlines	382	-0.025	0.000	0.001
easyJet	382	-0.229	-0.001	0.002
Non-bidding firms	0			

9 Concluding remarks

This paper assesses the welfare impacts of introducing competition in allocating scarce airport take-off and landing slots and provides handles to transition to a policy of periodic slot auctions. Methodologically, a new approach is taken to understand the effects of introducing an auction in a setting without relying on value-identifying bids. To this end, a structural flight-level model is developed and estimated, and slot values are identified in equilibrium as the incremental profit that firms generate by utilizing specific slots. Certain elements are introduced in the model to capture the essence of what makes airport slots valuable, including factors that make flights at favorable departure times and with short layovers more attractive for consumers. The model's accuracy is demonstrated through its close alignment with industry data, including metrics like gross profit margins from Air France, average slot prices at Heathrow, and the behavior of carriers seeking additional slots after the 2019 Aigle Azul bankruptcy.

The estimation results indicate significant network effects in costs and demand, influencing the slot values and favoring larger carriers in a market-based slot allocation. Counterfactual policy simulations demonstrate that both participating firms and consumers benefit from the auction. An interesting outcome of the estimated consumer preference for convenience is their overall preference for a mechanism without quantity caps or set-asides. In simpler terms, consumers prioritize the convenience benefits offered by larger dominant carriers over the advantages of increased competition from smaller network firms, at least within the context of the examined auctions. These findings from the model shed light on the central ambiguity in auctions with aftermarkets, where the welfare implications of efficient mechanisms are uncertain.

Regarding taxpayer gains, the projected proceeds of approximately €2.2 billion from auctioning 10% of slots in Paris align closely with those from the most recent French auction for radio spectrum. As the overall welfare gains of the policy are positive, firms can be made at least indifferent by restituting some of the auction proceeds via reduced airport fees, aiding the political feasibility of transitioning to a market-based system.

The analysis is done independently of the exact mechanism to be put in place and leans on the existence of an efficient mechanism to allocate slots based on the estimated slot values. Slot values are computed separately, considering only the complementarity with the carrier's existing flight network, and can be considered an approximation of the outcome of more complicated ascending auction designs. This is supported by the limited complementarity observed in the (already packaged) WSPs, at least for larger carriers. Moreover, the overall network size of carriers is only minimally impacted by the auction. With the caveat that this is based on the estimated stand-alone slot values, this limits the concern that dramatically changing network sizes would increasingly affect the slot values throughout the auction.

The estimation results serve as the basis for further study into the ideal auction design and to motivate why this would benefit society. An ideal policy would likely involve promoting environmental sustainability of the sector and other non-standard objectives that are are the forefront of the policy debate.

References

Abrevaya, Jason and Stephen G Donald. A gmm approach for dealing with missing data on regressors. *Review of Economics and Statistics*, 99(4):657–662, 2017.

Aguirregabiria, Victor and Chun-Yu Ho. A dynamic oligopoly game of the us airline industry: Estimation and policy experiments. *Journal of Econometrics*, 168(1):156–173, 2012.

Andrews, Donald WK and Biao Lu. Consistent model and moment selection procedures for gmm estimation with application to dynamic panel data models. *Journal of econometrics*, 101(1):123–164, 2001.

Armstrong, Timothy B. Large market asymptotics for differentiated product demand estimators with economic models of supply. *Econometrica*, 84(5):1961–1980, 2016.

Aryal, Gaurab, Charles Murry, and Jonathan W Williams. Price discrimination in international airline markets. *arXiv preprint arXiv:2102.05751*, 2021.

Athey, Susan, Dominic Coey, and Jonathan Levin. Set-asides and subsidies in auctions. *American economic journal: microeconomics*, 5(1):1–27, 2013.

Ausubel, Lawrence M, Peter Cramton, and Paul Milgrom. The clock-proxy auction: A practical combinatorial auction design. *Handbook of spectrum auction design*, pages 120–140, 2006.

Bajari, Patrick and Ali Hortaçsu. The winner's curse, reserve prices, and endogenous entry: Empirical insights from ebay auctions. *RAND Journal of Economics*, pages 329–355, 2003.

Ball, Michael O, Lawrence M. Ausubel, Frank Berardino, Peter Cramton, George Donohue, Mark Hansen, and Karla Hoffman. Market-based alternatives for managing congestion at new york's laguardia airport. AirNeth Annual Conference, The Hague, 2007.

Berry, Steven and Panle Jia. Tracing the woes: An empirical analysis of the airline industry. *American Economic Journal: Microeconomics*, 2(3):1–43, 2010.

Berry, Steven, Michael Carnall, and Pablo T Spiller. Airline hubs: costs, markups and the implications of customer heterogeneity. In Lee, Darin, editor, *Competition policy and antitrust*, pages 183–213. Elsevier, 2006.

Berry, Steven T. Airport presence as product differentiation. *The American Economic Review*, 80(2):394–399, 1990.

Berry, Steven T. Estimation of a model of entry in the airline industry. *Econometrica: Journal of the Econometric Society*, pages 889–917, 1992.

Berry, Steven T. Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, pages 242–262, 1994.

Berry, Steven T and Philip A Haile. Identification in differentiated products markets using market level data. *Econometrica*, 82(5):1749–1797, 2014.

Berry, Steven T., James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, 1995.

Bontemps, Christian, Cristina Gualdani, and Kevin Remmy. Price competition and endogenous product choice in networks: Evidence from the us airline industry. *Toulouse School of Economics*, 2021.

Borenstein, Severin. On the efficiency of competitive markets for operating licenses. *The Quarterly Journal of Economics*, 103(2):357–385, 1988.

Borenstein, Severin. Hubs and high fares: dominance and market power in the us airline industry. *The RAND Journal of Economics*, pages 344–365, 1989.

Borenstein, Severin and Janet Netz. Why do all the flights leave at 8 am?: Competition and departure-time differentiation in airline markets. *International journal of industrial organization*, 17(5):611–640, 1999.

Brenkers, Randy and Frank Verboven. Liberalizing a distribution system: the european car market. *Journal of the European Economic Association*, 4(1):216–251, 2006.

Brueckner, Jan K and Pablo T Spiller. Economies of traffic density in the deregulated airline industry. *The Journal of Law and Economics*, 37(2):379–415, 1994.

Brueckner, Jan K, Nichola J Dyer, and Pablo T Spiller. Fare determination in airline hub-and-spoke networks. *The RAND Journal of Economics*, pages 309–333, 1992.

Brueckner, Jan K, Matthew E Kahn, and Jerry Nickelsburg. How do airlines cut fuel usage, reducing their carbon emissions? 2023.

Cantillon, Estelle and Martin Pesendorfer. Combination bidding in multi-unit auctions. 2007.

Cardell, N Scott. Variance components structures for the extreme-value and logistic distributions with application to models of heterogeneity. *Econometric Theory*, 13(2):185–213, 1997.

Ciliberto, Federico and Jonathan W Williams. Limited access to airport facilities and market power in the airline industry. *The Journal of Law and Economics*, 53(3):467–495, 2010.

Ciliberto, Federico and Jonathan W Williams. Does multimarket contact facilitate tacit collusion? inference on conduct parameters in the airline industry. *The RAND Journal of Economics*, 45(4):764–791, 2014.

Ciliberto, Federico, Charles Murry, and Elie Tamer. Market structure and competition in airline markets. *Journal of Political Economy*, 129(11):2995–3038, 2021.

Conlon, Christopher and Jeff Gortmaker. Best practices for differentiated products demand estimation with pyblp. *The RAND Journal of Economics*, 51(4):1108–1161, 2020.

Cramton, Peter, Evan Kwerel, Gregory Rosston, and Andrzej Skrzypacz. Using spectrum auctions to enhance competition in wireless services. *The Journal of Law and Economics*, 54(S4):S167–S188, 2011.

Dagenais, Marcel G. The use of incomplete observations in multiple regression analysis: A generalized least squares approach. *Journal of Econometrics*, 1(4):317–328, 1973.

Dardanoni, Valentino, Salvatore Modica, and Franco Peracchi. Regression with imputed covariates: A generalized missing-indicator approach. *Journal of Econometrics*, 162(2):362–368, 2011.

Day, B and Paul Milgrom. Optimal incentives in core-selecting auctions. *Handbook of Market Design*, 2010.

De Wit, Jaap and Guillaume Burghouwt. Slot allocation and use at hub airports, perspectives for secondary trading. *European Journal of Transport and Infrastructure Research*, 8(2), 2008.

D'Haultfœuille, Xavier, Ao Wang, Philippe Février, and Lionel Wilner. Estimating the gains (and losses) of revenue management. *arXiv preprint arXiv:2206.04424*, 2022.

Doraszelski, Ulrich, Katja Seim, Michael Sinkinson, and Peichun Wang. Ownership concentration and strategic supply reduction. Technical report, National Bureau of Economic Research, 2017.

Einav, Liran and Jonathan Levin. Empirical industrial organization: A progress report. *Journal of Economic Perspectives*, 24(2):145–162, 2010.

Eizenberg, Alon. Upstream innovation and product variety in the us home pc market. *Review of Economic Studies*, 81(3):1003–1045, 2014.

Ershov, Daniel and David J Salant. Auctions and mobile market competition: Evidence from european 4g auctions. *Available at SSRN 4168166*, 2022.

Fan, Ying and Chenyu Yang. Competition, product proliferation, and welfare: A study of the us smartphone market. *American Economic Journal: Microeconomics*, 12(2):99–134, 2020.

Fox, Jeremy T and Patrick Bajari. Measuring the efficiency of an fcc spectrum auction. *American Economic Journal: Microeconomics*, 5(1):100–146, 2013.

Fu, Xiaowen, Tae Hoon Oum, Ruowei Chen, and Zheng Lei. Dominant carrier performance and international liberalization–the case of northeast asia. *Transport Policy*, 43:61–75, 2015.

Fukui, Hideki. Do carriers abuse the slot system to inhibit airport capacity usage? evidence from the us experience. *Journal of Air Transport Management*, 24:1–6, 2012.

Gayle, Philip G. On the efficiency of codeshare contracts between airlines: is double marginalization eliminated? *American Economic Journal: Microeconomics*, 5(4):244–273, 2013.

Gentry, Matthew, Tatiana Komarova, and Pasquale Schiraldi. Preferences and performance in simultaneous first-price auctions: A structural analysis. *The Review of Economic Studies*, 90(2):852–878, 2023.

Goldberg, Pinelopi Koujianou and Frank Verboven. The evolution of price dispersion in the european car market. *The Review of Economic Studies*, 68(4):811–848, 2001.

Gourieroux, Christian and Alain Monfort. On the problem of missing data in linear models. *The Review of Economic Studies*, 48(4):579–586, 1981.

Grether, David M, R Mark Isaac, and Charles R Plott. Alternative methods of allocating airport slots: Performance and evaluation. 1979.

Gretschko, Vitali, Martin Bichler, and Maarten Janssen. Bargaining in spectrum auctions: A review of the german auction in 2015. 2016.

Grigolon, Laura and Frank Verboven. Nested logit or random coefficients logit? a comparison of alternative discrete choice models of product differentiation. *Review of Economics and Statistics*, 96(5):916–935, 2014.

Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong. Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68(3):525–574, 2000.

Hansen, Lars Peter. Large sample properties of generalized method of moments estimators. *Econometrica*, pages 1029–1054, 1982.

Hazlett, Thomas W and Roberto E Muñoz. A welfare analysis of spectrum allocation policies. *The RAND Journal of Economics*, 40(3):424–454, 2009.

Ivaldi, Marc, Milena Petrova, and Miguel Urdanoz. Airline cooperation effects on airfare distribution: An auction-model-based approach. *Transport Policy*, 115:239–250, 2022.

Janssen, Maarten CW and Vladimir A Karamychev. Auctions, aftermarket competition, and risk attitudes. *International Journal of Industrial Organization*, 27(2):274–285, 2009.

Jehiel, Philippe and Benny Moldovanu. Efficient design with interdependent valuations. *Econometrica*, 69(5): 1237–1259, 2001.

Kasberger, Bernhard. When can auctions maximize post-auction welfare? *Available at SSRN 3519866*, 2020.

Kim, Sang Won, Marcelo Olivares, and Gabriel Y Weintraub. Measuring the performance of large-scale combinatorial auctions: A structural estimation approach. *Management Science*, 60(5):1180–1201, 2014.

Kong, Yunmi. Sequential auctions with synergy and affiliation across auctions. *Journal of Political Economy*, 129(1):148–181, 2021.

Kösters, Till, Marlena Meier, and Gernot Sieg. Effects of the use-it-or-lose-it rule on airline strategy and climate. *Transportation Research Part D: Transport and Environment*, 115:103570, 2023.

Krishna, Vijay. *Auction theory.* Academic press, 2009.

Laffont, Jean-Jacques, Herve Ossard, and Quang Vuong. Econometrics of first-price auctions. *Econometrica: Journal of the Econometric Society*, pages 953–980, 1995.

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

Levine, Michael E. Airline competition in deregulated markets: theory, firm strategy, and public policy. *Yale J. on Reg.*, 4:393, 1986.

Milgrom, Paul. Putting auction theory to work: The simultaneous ascending auction. *Journal of political economy*, 108(2):245–272, 2000.

Miller, Nathan, Steven Berry, Fiona Scott Morton, Jonathan Baker, Timothy Bresnahan, Martin Gaynor, Richard Gilbert, George Hay, Ginger Jin, Bruce Kobayashi, and others. On the misuse of regressions of price on the hhi in merger review. *Journal of Antitrust Enforcement*, 10(2):248–259, 2022.

Miller, Nathan H and Matthew C Weinberg. Understanding the price effects of the millercoors joint venture. *Econometrica*, 85(6):1763–1791, 2017.

Miranda, Victor AP and Alessandro VM Oliveira. Airport slots and the internalization of congestion by airlines: An empirical model of integrated flight disruption management in brazil. *Transportation Research Part A: Policy and Practice*, 116:201–219, 2018.

Miravete, Eugenio J, Katja Seim, and Jeff Thurk. Market power and the laffer curve. *Econometrica*, 86(5): 1651–1687, 2018.

Odoni, Amedeo R. A review of certain aspects of the slot allocation process at level 3 airports under regulation 95/93. *MIT International Center for Air Transportation Report No. ICAT-2020-09*, 2020.

Paarsch, Harry J. Deciding between the common and private value paradigms in empirical models of auctions. *Journal of econometrics*, 51(1-2):191–215, 1992.

Pakes, Ariel. Alternative models for moment inequalities. *Econometrica*, 78(6):1783–1822, 2010.

Pakes, Ariel, Jack Porter, Kate Ho, and Joy Ishii. Moment inequalities and their application. *Econometrica*, 83 (1):315–334, 2015.

Park, Yongjoon. Structural remedies in network industries: An assessment of slot divestitures in the american airlines/us airways merger. *SSRN*, 2020.

Parkes, David C and Lyle H Ungar. Iterative combinatorial auctions: Theory and practice. *Aaai/iaai*, 7481:53, 2000.

Rassenti, Stephen J, Vernon L Smith, and Robert L Bulfin. A combinatorial auction mechanism for airport time slot allocation. *The Bell Journal of Economics*, pages 402–417, 1982.

Reitzes, James D, Brendan McVeigh, Nicholas Powers, and Samuel Moy. Competitive effects of exchanges or sales of airport landing slots. *Review of Industrial Organization*, 46:95–125, 2015.

Ribeiro, Nuno Antunes, Alexandre Jacquillat, António Pais Antunes, Amedeo R Odoni, and João P Pita. An optimization approach for airport slot allocation under iata guidelines. *Transportation Research Part B: Methodological*, 112:132–156, 2018.

Rubin, Donald B. Inference and missing data. *Biometrika*, 63(3):581–592, 1976.

Seymour, Kyle, Maximilian Held, Gil Georges, and Konstantinos Boulouchos. Fuel estimation in air transportation: Modeling global fuel consumption for commercial aviation. *Transportation Research Part D: Transport and Environment*, 88:102528, 2020.

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

Wollmann, Thomas G. Trucks without bailouts: Equilibrium product characteristics for commercial vehicles. *American Economic Review*, 108(6):1364–1406, 2018.

Xiao, Mo and Zhe Yuan. License complementarity and package bidding: Us spectrum auctions. *American Economic Journal: Microeconomics*, 14(4):420–64, 2022.

Online Appendices

A Fare sampling process

This appendix details the DGAC observed fare sampling process in relation to the treatment of missing values in the survey-weighted lowest offered fare of the flight (P_i).

This section specifies the demand side as a simple linear regression model to clarify the relationship between the process generating missing values in the price variable and the outcomes of interest. Generic subscripts for observation i (which represent a consumer-flight combination) are used, assuming that the relevant population can be accurately described with one consumer type and that there is no correlation of preferences requiring a nested error structure. Dependent variable Y_i represents indirect consumer utility from taking flight i, P_i the average fare, X_i' a row vector of observed flight characteristics that are never missing (including intercept), ξ_i the unobservable flight quality, ϵ_i a scalar unobservable assumed i.i.d. across individuals and flights, and $\theta = \{\alpha, \beta\}$ parameters to estimate with a focus on the disutility of price (α):

$$Y_i = \alpha P_i + X_i'\beta + \xi_i + \epsilon_i \tag{A.1}$$

with $\mathbb{E}[X_i\epsilon_i] = \mathbb{E}[X_i\xi_i] = \mathbb{E}[P_i\epsilon_i] = 0$ and $\mathbb{E}[P_i\xi_i] \neq 0$. Furthermore, to describe the event that the price is missing, define M_i as:

$$M_i = \begin{cases} 1 \text{ if } P_i \text{ is missing} \\ 0 \text{ if } P_i \text{ is not missing} \end{cases}$$

Note that the analysis in the paper is done at the flight level requiring the usual logit transformation of (A.1) into market shares (e.g., adopting a Gumbell distribution for ϵ, specifying the outside option and consumer choice problem). The methods to deal with the missingness of fare data that are discussed below also apply when P_i is endogenous –conditional on the availability of set of valid instruments (row vector) $Z_i' = [\{Z_i^k\}_{k=1,\dots,K}]$ for $K \geq 1$, and (when necessary) conditioning on these as well with $W_i = [X_i', Z_i']'$. In the spirit of Rubin (1976), examining what drives M_i is crucial.

Flight-level prices P_i are the survey-weighted average ticket-level offered prices collected periodically by the DGAC web-scraping algorithm. The algorithm works as follows. Every other day, it collects ticket prices for a return flight offered for a predetermined set of "products" and "profiles", which amount to the combination of departure airport, destination airport, ticketing carrier, number of days of stay (3, 7, or 21), in how many days the flight takes place (2, 11, 20, 44, or 90), booking class (economy or business), with flexible booking conditions or not, and whether it includes a stopover or not.[57] For each combination of these search characteristics $\mathcal{A}$, the robot records the lowest offered price if there are multiple.[58] After dropping flights with a stopover (due to not having quantity data for such flights), the average fare for flight i is obtained as the survey-weighted average of

[57]To give an example, the algorithm loads prices for a return flight from CDG to BCN, without connection, for a departure date 2 days ahead and a return date 5 days ahead, in economy class, with flexible booking conditions, and with Air France.

[58]The algorithm thus mimics consumers selecting the lowest available fare for a given itinerary, as in the dynamic pricing model of Williams (2022), with the addition of also selecting the carrier, return leg, class, flexibility of booking conditions, and connection.

recorded prices related to flight i ($\mathcal{P}^s_{ai}$):[59]

$$P_i = \frac{1}{S_i}\sum_{s=1}^{S_i}\mathcal{P}^s_{ai}\phi_a \tag{A.2}$$

Survey weights ϕ_a are provided by civil aviation authority DGAC to render the recorded prices representative of air travel spending by the relevant population according to passenger surveys.[60] S_i (the total number of times a price for flight i is recorded) varies by flight due to i) $\mathcal{A}$ not containing the departure time and ii) unavailability of tickets when the departure day approaches. Importantly, the fact that the robot does not specifically condition on the departure time explains why the distribution of departure hours differs by M_i. Flights departing at unfavorable hours (to consumers) have lower equilibrium fares and are therefore *oversampled* by the algorithm.

With the above sampling mechanism in mind, it is easy to show that M_i does not depend on Y_i conditional on observables under reasonable assumptions. As standard in structural models of oligopoly markets, prices are conceived to be a function of demand characteristics (X_i, affecting mark-ups), a flight-level scalar unobservable "quality" (ξ_i, the root cause of endogeneity in these models), and cost. The latter are referred to as cost shifters (Z_i) when applied to address price endogeneity using a GMM model, based on the identifying assumption that $\mathbb{E}[\xi_i W_i] = 0$ for $W_i' = [X_i', Z_i']'$ as above and with instruments Z_i excluded from consumer utility. I next use these conventions to explicitly model the set of offered prices during simulation s of the price scraping algorithm for search profile a. Offered prices $\mathcal{P}^s_{ain}$ are also subscripted by i to indicate that the price is associated with (additional) characteristics for flight i (X_i', Z_i', ξ_i) and by n to denote one of the offers if there are multiple. Hence $\mathcal{P}^s_{ain}$ can be specified as:

$$\mathcal{P}^s_{ain} = g(a, X_i, Z_i) + h(\xi_i) + \nu^s_{an}, \text{ with } \nu^s_{an} \perp (\xi_i, \epsilon_i) \tag{A.3}$$

for generic functions $g(.)$ and $h(.)$. The important restriction is that ξ_i is a proper flight-level unobservable, constant across multiple offers within a simulation, and independent of any offer-specific unobserved heterogeneity ν^s_{an}. Rather than selecting the minimum offer price across n conditional on (a, X_i, Z_i) in simulation s, the algorithm selects price:

$$\begin{aligned}\mathcal{P}^s_{ai} &= \min(\mathcal{P}^s_{ain}|\mathcal{A} = a) \\ &= \min(g(a, X_i, Z_i) + h(\xi_i) + \nu^s_{an}) \\ &= \min(\nu^s_{an}) + g(a, X_i, Z_i) + h(\xi_i)\end{aligned} \tag{A.4}$$

with the last equality holding in particular due to $h(\xi_i)$ being constant across n. This underscores the earlier comment that unattractive departure times in X_i are undersampled as they are associated with higher prices in equilibrium (higher $g(a, X_i, Z_i)$). Also flights with higher cost, i.e. due to the use of a less fuel efficient aircraft, or with higher offer-specific idiosyncratic error are undersampled.

Importantly though, the ignorability condition of Dardanoni et al. (2011) holds in this context.[61] The

[59] Typically, $\mathcal{P}^s_{ai}$ is one of many recorded prices that match on the market-product definition (departure month, departure airport, destination airport, ticketing carrier, departure day, departure hour) regardless of the number of days of the stay, days ahead the flight takes place, booking class, and flexibility of booking conditions.

[60] DGAC produce survey-weighted ticket prices as part of the travel component of the French CPI, and have kindly generated survey weights ϕ_a for this study to match expenditure by travellers departing from CDG or ORY airports only.

[61] Dardanoni et al. (2011) use the term *ignorability* for when M_i and (Y_i, P_i) are conditionally independent given W_i. This renders both the complete case and linear imputation method (e.g., Dagenais (1973), Gourieroux and Monfort (1981)) consistent, with the former being potentially less precise and the latter potentially introducing bias through the imputation equation.

equation below formalizes this by letting $DepHours(.)$ be the set of departure hours associated with selected minimum prices $\mathcal{P}^s_{ai}$, $a(i)$ the set of product-market characteristics associated with observation i (e.g. departure airport, arrival airport, ticketing carrier, month), and $S_a(i)$ the number of simulations done for these characteristics. P_i is only missing if:

$$DepHour_i \notin DepHours\left(\{\mathcal{P}^s_{ai}\}_{a=a(i),s=1,\dots,S_a(i)}\right) \tag{A.5}$$

so that $M_i \perp Y_i | (X_i, Z_i)$ under the restrictions in (A.3). Under these restrictions, imputation methods are more efficient than the complete case estimator, consistent and asymptotically normal (Dagenais (1973), Gourieroux and Monfort (1981), Abrevaya and Donald (2017)). This motivates the imputation approach the paper takes to deal with missing prices. Imputing missing fares is merely a prediction problem; no causal relationship is sought after at this stage. To determine which combination of variables best predicts the price of a flight, various (machine learning) methods are explored alongside more standard regressions (see section B).

Another remark here is that a more stringent assumption that prices are Missing Completely At Random (MCAR) does not hold. MCAR assumes that M_i is statistically independent of (P_i, W_i, ϵ_i) and any unobserved price heterogeneity conditional on W_i. When MCAR holds, this justifies discarding observations with $M_i = 1$ and adopting the complete case estimator –as is done in most empirical studies dealing with missing data published in the top econ journals.[62] However, MCAR is violated as the distribution of departure hours differs by M_i (see figure A1). To see why this is the case, recall that the algorithm collects prices for different ticket characteristics for each combination of departure day of the week, origin, destination, and carrier. However, it does not specifically search for flights departing at a certain hour or time. As a result, even the reweighed data might not fully reflect consumer preferences for slots. The lowest price is recorded if multiple departure times are offered within the (narrow) ticket search.

B Machine Learning Imputation Methods

This appendix describes the application of machine learning techniques alongside conventional regression models to impute missing fares.

Three economically-motivated regression models that relate the fare to variables in X and markup- and cost-shifters are specified based on what we know about the product/industry. In addition, three machine learning methods are considered: Random Forest, Ranger, and Ridge Regression. Random Forest and Ranger are based on partitioning the data into increasingly fine *bins* with a bin-specific (average) fare. These algorithms search over the set of variables (and values of factor variables) for the best-fitting partitioning. At the same time, Ridge Regression takes all available values and gives them the best-fitting weights (which may include 0). The 48,241 observations with non-missing Fares (40.0% of the full dataset) are split 80-20 into training and test sets. This operation results in a training set of 38,593 observations. For the expert-dictated or economically motivated models, the coefficients are estimated by OLS on the full training dataset. The machine learning methods split the training set into one half where the model is trained and one half where the model is validated. The

[62] Abrevaya and Donald (2017) review all empirical papers published in the American Economic Review, Quarterly Journal of Economics, the Journal of Human Resources, and the Journal of Labor Economics published between 2006 and 2008 and find that 40 percent deal with missing data. In roughly 70 percent of those articles observations with missing values are dropped (referred to as the complete data method), and the remainder use linear imputation of missing values and/or a dummy variable (or missing-indicator) method.

(a) By day (b) All days

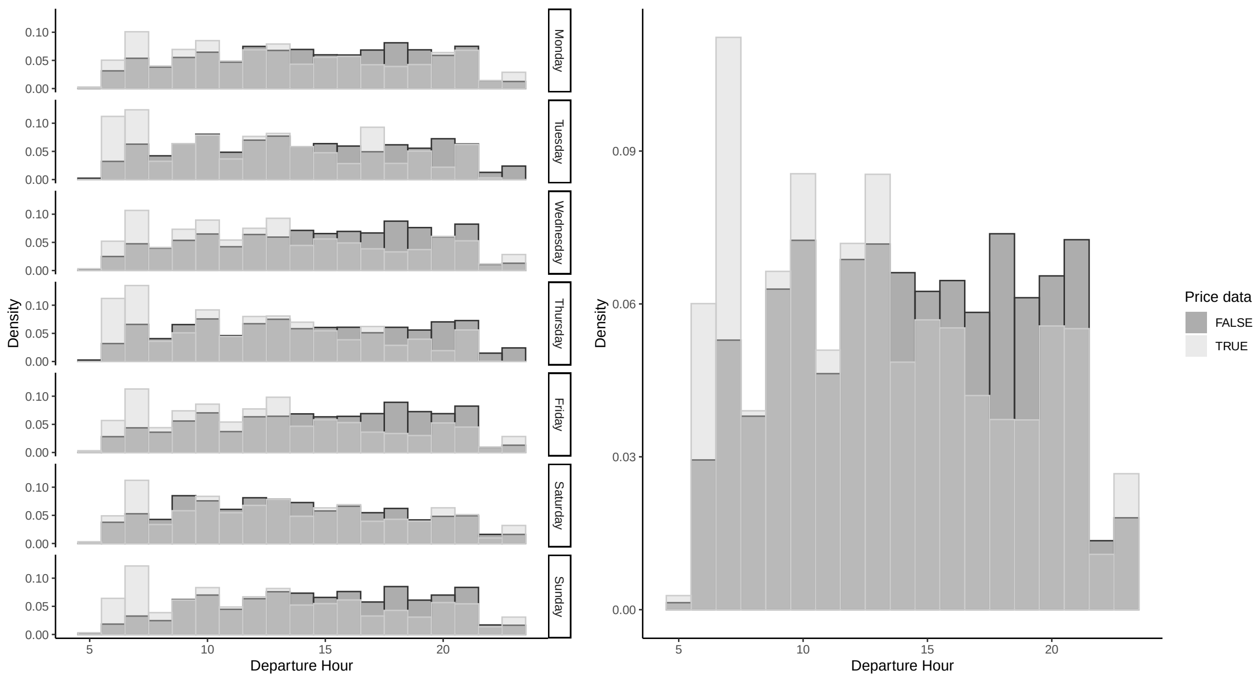

Figure A1: PMF of departure hour, separately for whether or not there is price data available for flight

remaining 9,648 observations constitute the test set, which always leaves aside any estimation or training, to avoid overfitting. Fares are predicted based on the trained/estimated models and compared to the true values in the test set, and the fit is summarised by the RMSE and R^2 in the test dataset.

The set of potential predictor variables include: X and W reported in the paper, BLP instruments based on X, additional instruments included in the analysis (sum of flights departing within a 2-hour window on the same day, sum of seats for those flights, both for own and rival firms, average fuel price per seat based on the price per dollar of Kerosene grade Jet fuel and the dollar/euro exchange rate in different periods before the flight), the market share, the market share^2, the market size, destination-specific variables such as the absolute deviation in temperature with Paris, and other flight-level observables not included in X and W. Some dataset pre-processing is performed to avoid multicollinearity (excluding variables with a correlation coefficient larger than 0.9), and the final training sample contains 104 variables. The machine learning models are trained using the `caret` package in `R`. The Random Forest, Ranger, and Ridge Regression methods are first trained on only 5,000 observations to limit the computation time, and the best-performing method is then also trained on the full training set of 38,593 observations. The performance of these seven models is summarized in table B1.

The first thing to note is that all simple OLS expert-dictated models do very well in predicting fares. Model B has the lowest R^2 and the fewest outliers among the OLS-based models (table B2). Ridge regression does better than the best OLS model even on the smaller training set of 5,000 observations, improving the R^2 by 6 percentage points. However, as the imputation is based on continuous variables and weights, it results in 0.2% $\hat{P}_{jt} < 0$ –similar to the OLS-based models. Random Forest and Ranger do not have that issue, and they result in an even better fit. The Full Ranger model, estimated on the full training dataset, improves the R^2 to an impressive 0.9 and lowers the RMSE to 81.71. It can be concluded that this machine learning method predicts with very high accuracy the average flight-level fares that are set based on the available data, and the structural demand estimation in the paper is based on fares imputed with this model (if missing).

Table B1: Performance imputation models: Fit

	RMSE	R^2	No. obs. Train	Train Control Method	CV number	CV repeated	User time (hr)
Model A	1.38	0.74	38838	NaN			$< 10s$
Model B	1.42	0.73	38838	NaN			$< 10s$
Model C	1.37	0.75	38838	NaN			$< 10s$
Random Forest	1.09	0.84	5000	repeatedcv	10.00	10.00	0.18774138888889
Ranger	1.09	0.84	5000	repeatedcv	10.00	10.00	0.0243886111111109
Ridge	1.32	0.77	5000	repeatedcv	10.00	10.00	0.00062249999999848
Full Ranger	0.91	0.89	38838	repeatedcv	10.00	10.00	11.0538288888889

Notes. Notes. Performance of the seven different regression/machine learning models. Control method "repeated CV" stands for repeated Cross Validation, and # CV indicates the number of times the model is cross-validated, each based on a 10-fold split of the Training sample. This means that the Full Ranger model is fitted 100 times on the Training sample, selecting the best-performing partitioning among those, and the performance of this model is evaluated in the training sample.

Table B2: Details imputation: distribution and outliers

	Nobs	Min	1st Q.	Median	Mean	3rd Q.	Max	Outlier (%)
Non-missing Fares	48793	Min. : 0.39	1st Qu.: 1.23	Median : 1.98	Mean : 3.22	3rd Qu.: 4.29	Max. :19.27	
Model A	120053	Min. :-0.6863	1st Qu.: 1.2624	Median : 1.9035	Mean : 2.7684	3rd Qu.: 3.1522	Max. :19.2665	1.1612
Model B	120053	Min. :-0.7189	1st Qu.: 1.2645	Median : 1.9080	Mean : 2.7651	3rd Qu.: 3.1292	Max. :19.2665	1.082
Model C	120053	Min. :-0.6711	1st Qu.: 1.2740	Median : 1.9376	Mean : 2.7848	3rd Qu.: 3.1659	Max. :19.2665	1.1503
Random Forest	120053	Min. : 0.3895	1st Qu.: 1.3989	Median : 2.0366	Mean : 2.9275	3rd Qu.: 3.0343	Max. :19.2665	0
Ranger	120053	Min. : 0.3895	1st Qu.: 1.4009	Median : 2.0374	Mean : 2.9258	3rd Qu.: 3.0127	Max. :19.2665	0
Ridge	120053	Min. : 0.2184	1st Qu.: 1.3556	Median : 1.8624	Mean : 2.8538	3rd Qu.: 3.0154	Max. :19.2665	0.0117
Ranger Full	120053	Min. : 0.3895	1st Qu.: 1.3444	Median : 2.0699	Mean : 3.0235	3rd Qu.: 3.5102	Max. :19.2665	0

Notes. Notes. Quantiles and statistics of the predicted $\hat{P}_{jt}$ based on the relevant regression/machine learning model. An outlier is defined as $\hat{P}_{jt} < 0$ or $\hat{P}_{jt} > 1204.13$ (the sample maximum of observed P_{jt}).

It is also interesting to learn from the machine which variables are the most important (statistically) predictor of prices in this market. Table B3 lists the 50 most important variables for the four models and their scaled importance within the model. Whether the destination requires a LongHaul flight (a combination of Europe dummy and Distance category) is by far the most important one in the Ranger Full model, followed by the CDG departure airport dummy, and aircraft-distance specific Fuel Consumption variables.

C Data: additional descriptive statistics and results

This appendix contains additional descriptive statistics and results referenced in the text.

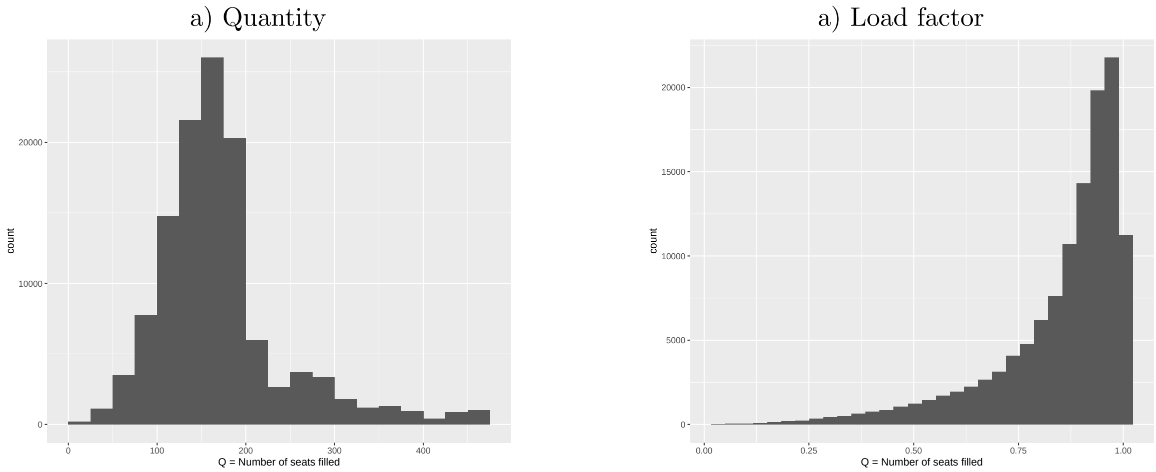

Figure C1: Quantity and load factor (quantity/capacity) across flights

Table B3: Machine learning models: 50 most important predictors of P_{jt}

Var. name	Random Forest	Ranger	Ridge Regression	Ranger Full
Distance	100.00	100.00	100.00	100.00
Distance2	91.43	99.59	2.26	87.04
outsideshare	14.67	20.00	0.86	21.98
UniqueFaresFlight	9.20	18.48	0.37	20.58
ZSpacingSumOwn	3.00	8.48	0.33	9.86
SHbusinessClass	2.76	7.60	0.33	8.64
log_FuelDistperSeatL9	2.72	6.40	0.27	6.79
FuelDistperSeatL9	2.49	6.11	0.23	6.68
FuelBurnperSeatL9	2.41	5.82	0.22	4.99
LoadFactor	2.18	5.80	0.22	4.94
DepartureDayMONRushHour	2.08	5.07	0.20	4.33
log_FuelBurnperSeatL9	2.05	4.30	0.18	4.30
PotConn6h_DEP	1.93	3.68	0.17	4.12
PotConn6h_ARR	1.87	3.62	0.17	3.30
RushHourMorningExp	1.74	3.54	0.13	3.11
CDG	1.55	3.42	0.13	3.05
Quantity	1.53	3.23	0.12	3.02
FlagCarrier	1.48	3.01	0.12	3.01
ZSpacingDestSumOwn	1.46	2.99	0.12	3.00
Exceed90thSpacing	1.43	2.95	0.12	2.95
ZSpacingDestSumOther	1.41	2.80	0.11	2.90
ZWeekdayRushMorningSumOwn	1.32	2.76	0.11	2.68
RushHourEveningExp	1.27	2.74	0.07	2.62
DepartureDateDay	1.27	2.54	0.07	2.61
k	1.19	2.49	0.07	2.58
Herfendahl	1.18	2.41	0.05	2.57
Destinations	1.15	2.37	0.05	2.32
HerfendahlDest	1.12	2.35	0.05	2.27
ZWeekdayRushEveningSumOwn	1.09	2.27	0.05	2.17
ZPotConn6h_ARRSumOwn	1.08	2.08	0.05	2.12
MediumHaul	1.08	2.05	0.05	2.11
ZPotConn6h_DEPSumOwn	1.06	2.01	0.04	2.03
Spacing	1.06	1.96	0.04	1.93
DepartureDayTUERushHour	1.01	1.94	0.04	1.83
RushHour	1.00	1.92	0.03	1.81
ZWeekdayRushMorningSumOther	1.00	1.88	0.03	1.80
prodshare	0.99	1.83	0.03	1.79
Month	0.97	1.69	0.02	1.78
marketshare2	0.95	1.69	0.02	1.77
ZFlightsSumOwn	0.95	1.66	0.02	1.75
marketshare	0.92	1.62	0.01	1.72
ZFlightsDestSumOwn	0.91	1.62	0.01	1.71
HubDest	0.89	1.57	0.01	1.70
FuelBurn	0.83	1.53	0.01	1.67
ZSeatsFlightsSumOther	0.80	1.52	0.01	1.66
Destinations2	0.80	1.52	0.01	1.53
SpacingDest	0.72	1.48	0.01	1.52
ArrivalHour	0.71	1.39	0.01	1.50
ZSHfirstClassSumOwn	0.70	1.28	0.01	1.49
WeekdayRushHour	0.70	1.24	0.01	1.48

Notes. Notes. LongHaul is based on the Distance×Europe dummy for the destination airport, the variables ending with Lx are based on the JetFuel price Lagged by x weeks, Fuel*Dist*perSeat is based on a distance-based approximation of the Fuel*Burn*perSeat variables (which reflects the aircraft- and distance- specific fuel consumption per seat), variables ending with OWN are BLP-instruments summing over own flights, and ending with $RIVAL$ are based on rival flights, the Ridge regression contains almost only destination and carrier fixed effects (indicated by their IATA codes).

Table C1: Slot changes when perturbing network by up to 1 flight

Flights S19 (number)	Flights change S19-S18 (number)	Additional Slots (percent)	Removed Slots (percent)	Across Slot Shuffle (percent)	City	Country
434	-1	0	0	-0.23	San Francisco	United States
30	-1	0	0	-3.33	Djibouti	Djibouti
54	0	0	0	0	Figari	France
217	0	0	0	0	Osaka	Japan
217	0	0	0	0	Jeddah	Saudi Arabia
24	1	0	0	4.17	Calvi	France

Notes. Based on Global Schedules Data; selecting Air France departures from Paris Charles de Gaulle (CDG) airport. The table reports changes in flights between the summer season of 2018 (S18) and the summer season for 2019 (S19) for destinations where one flight is added (Calvi), one flight is removed (San Francisco and Djibouti), or where the number of flights is exactly the same in S18 as S19 (Figari, Osaka, and Jeddah), as indicated in the column *Nr Flights change S19*. The percent of flights in slots that is added to (or removed from) the network in this destination is given in the column *Additional Slots* (*Removed Slots*). The zeros in these columns indicate that the flights are added to existing slot windows, indicating that an extra week is added or removed from the slot series that Air France operated in S18. Positive numbers indicate that a new slot window is used. The column *Across Slot Shuffle* reports the average change in the number of flights scheduled in those slots as a percentage of the number of flights scheduled in S19.

Table C2: Flights added from Paris CDG or ORY in summer of 2019

	New Flights	City	Country	Carrier
New Destination, New Carrier	254	Calgary	Canada	WestJet
New Destination, New Carrier	161	Halifax	Canada	WestJet
New Destination, New Carrier	120	Fuzhou	China	Xiamen Airlines
Total new flights	535			
New Destination, Existing Carrier	150	Sao Paulo	Brazil	Aigle Azur
New Destination, Existing Carrier	30	Bari	Italy	Air France
New Destination, Existing Carrier	210	La Rochelle	France	Air France
New Destination, Existing Carrier	72	Quito	Ecuador	Air France
New Destination, Existing Carrier	210	Wroclaw	Poland	Air France
New Destination, Existing Carrier	60	Chongqing	China	Hainan Airlines
New Destination, Existing Carrier	60	Shenzhen	China	Hainan Airlines
New Destination, Existing Carrier	90	Vigo	Spain	Iberia
New Destination, Existing Carrier	30	Ivalo	Finland	Jetairfly
New Destination, Existing Carrier	50	Nador	Morocco	Transavia France
Total new flights	962			
Existing Destination, New Carrier	210	Dakar	Senegal	Air Senegal
Existing Destination, New Carrier	360	Vienna	Austria	Anisec
Existing Destination, New Carrier	117	Fort-de-France	Martinique	Openskies
Existing Destination, New Carrier	78	Montreal	Canada	Openskies
Existing Destination, New Carrier	133	New York	United States	Openskies
Existing Destination, New Carrier	136	Pointe A Pitre	Guadeloupe	Openskies
Existing Destination, New Carrier	120	Algiers	Algeria	Tassili Airlines
Total new flights	1154			
Total New Flights, New Destination or Carrier	2651			
Total New Flights, Existing Destinations and Carriers	32086			
Total Existing Destinations and Carriers	224			
Total New Destinations by New Carriers	4			
Total New Destinations by Existing Carriers	11			
Total Existing Destinations by New Carriers	8			

Notes. Based on Global Schedules Data; selecting departures from Paris Charles de Gaulle (CDG) or Paris Orly (ORY) airport. The table reports all flights offered in the summer season of 2019 (S19) that were not offered in the same season in 2018 (S18).

Table C3: Share of passengers to US destinations

Destination	DGAC	T100
ATL	0.10	0.11
BOS	0.06	0.07
EWR	0.06	0.10
IAD	0.11	0.08
JFK	0.27	0.27
LAX	0.11	0.14
MIA	0.08	0.06
ORD	0.08	0.06
SFO	0.12	0.10

Notes. Rows reflect the nine destinations that have commercial flights in both the international segment of the T-100 dataset and in the DGAC quantity data between April and September 2018, with the departure airport being CDG or ORY. The DOT T-100 (International segment) dataset is obtained from https://www.transtats.bts.gov/DL_SelectFields.aspx?gnoyr_VQ=FJE&QO_fu146_anzr=Nv4%20Pn44vr45. It is based on a 10% random sample of ticket sales for direct international flights originating or departing in the US, as reported by the airline carriers. The DGAC dataset contains the number of actual passengers for the universe of direct flights departing from CDG or ORY.

Table C4: Estimation Sample: Selected Carriers and Destinations

		Means (across destinations)	
N	In Sample	Carriers	Flights
119	TRUE	4.21	1227.77
285	FALSE	2.38	290.38

		Means (across carriers)	
N	In Sample	Dests.	Flights
57	TRUE	7.84	2931.14
63	FALSE	2.40	371.24

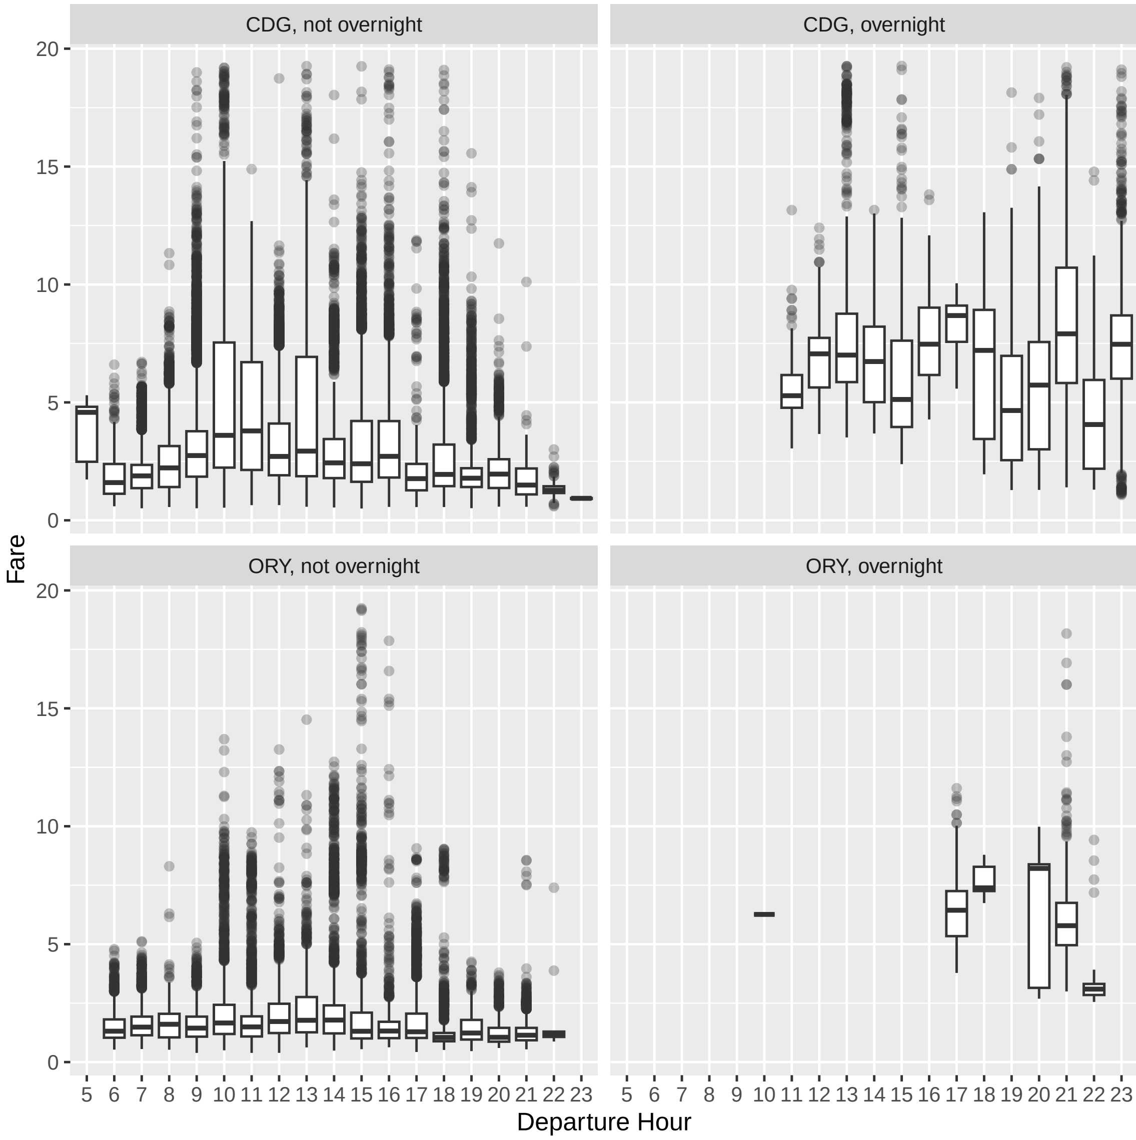

Figure C2: Price variation across departure hours, by airport and overnight flight indicator

Table C5: Carriers included in the sample

Carrier name	Number Destinations	Number Flights	Share of Quantity
Air France	170	88991	0.39
easyJet	50	14926	0.07
Transavia France	49	9867	0.05
Vueling Airlines	27	7466	0.04
Aigle Azur	13	2094	0.01
Delta Air Lines	13	2427	0.02
Air Algerie	10	2661	0.01
Corsair	10	1090	0.01
Air Caraibes	9	1040	0.01
Royal Air Maroc	8	2895	0.01
American Airlines	6	1226	0.01
Aegean Airlines	5	1057	0.00
Tunisair	5	1627	0.01
Air Corsica	4	2046	0.01
United Airlines	4	884	0.01
Air Transat	4	575	0.01
Turkish Airlines	3	1343	0.01
SAS	3	1393	0.01
Air Canada	3	547	0.01
Aeroflot	3	1510	0.01
Air China Limited	3	668	0.00
Air Austral	3	233	0.00
Lufthansa	2	2734	0.01
TAP	2	1640	0.01
Jet Airways	2	367	0.00
Aer Lingus	2	1076	0.00
China Eastern Airlines	2	525	0.00
Vietnam Airlines	2	308	0.00
Air Madagascar	2	53	0.00
Ukraine International	1	432	0.00
SWISS	1	1028	0.00
British Airways	1	1365	0.01
Finnair	1	1067	0.00
Iberia	1	1161	0.01
KLM	1	1164	0.01
Qatar Airways	1	640	0.01
CSA	1	645	0.00
Austrian	1	804	0.00
Middle East Airlines	1	432	0.00
El Al	1	560	0.00
Emirates	1	617	0.01
Singapore Airlines	1	216	0.00
TAROM	1	426	0.00
Cathay Pacific Airways	1	347	0.00
Thai Airways International	1	208	0.00
Atlasglobal	1	189	0.00
Air Mauritius	1	249	0.00
Asiana Airlines	1	154	0.00
Air Serbia	1	394	0.00
JAL	1	216	0.00
Korean Air Lines	1	233	0.00
Air Seychelles	1	13	0.00
Aeromexico	1	208	0.00
Etihad Airways	1	423	0.00
ANA	1	214	0.00
Air India	1	215	0.00
LATAM Airlines Brasil	1	186	0.00

Table C6: Flights in the sample with at most 10 passengers

CarrierName	Origin	arrivalCityName	Quantity	Capacity	DepartureDay	DepartureHour	DepartureDate
RoyalAirMaroc	ORY	Agadir	4	114	Thursday	22	2018-08-23
RoyalAirMaroc	ORY	Agadir	4	114	Friday	22	2018-08-24
RoyalAirMaroc	ORY	Agadir	8	114	Tuesday	22	2018-08-28
RoyalAirMaroc	ORY	Agadir	4	114	Wednesday	22	2018-08-29
AirCorsica	ORY	Ajaccio	6	180	Monday	22	2018-08-27
AigleAzur	ORY	Algiers	5	180	Sunday	7	2018-09-30
AirFrance	ORY	Bordeaux	6	145	Monday	7	2018-04-02
AirFrance	CDG	Bordeaux	6	145	Monday	12	2018-05-07
RoyalAirMaroc	CDG	Casablanca	5	114	Thursday	7	2018-08-23
RoyalAirMaroc	CDG	Casablanca	7	159	Saturday	7	2018-08-25
RoyalAirMaroc	CDG	Casablanca	7	114	Sunday	7	2018-08-26
RoyalAirMaroc	CDG	Casablanca	6	159	Tuesday	7	2018-08-28
RoyalAirMaroc	CDG	Casablanca	8	159	Wednesday	7	2018-08-29
RoyalAirMaroc	CDG	Casablanca	10	159	Thursday	7	2018-08-30
RoyalAirMaroc	CDG	Casablanca	8	159	Wednesday	7	2018-09-05
RoyalAirMaroc	CDG	Casablanca	9	159	Friday	7	2018-09-07
Tunisair	ORY	Djerba	9	148	Tuesday	19	2018-08-28
Tunisair	ORY	Djerba	6	148	Wednesday	19	2018-08-29
Tunisair	ORY	Djerba	9	148	Monday	19	2018-09-03
Tunisair	ORY	Djerba	8	148	Tuesday	19	2018-09-04
AirFrance	ORY	Marseille	5	180	Tuesday	17	2018-09-04
AirFrance	ORY	Toulouse	8	180	Tuesday	21	2018-04-03
AirFrance	CDG	Tunis	10	169	Friday	7	2018-08-31

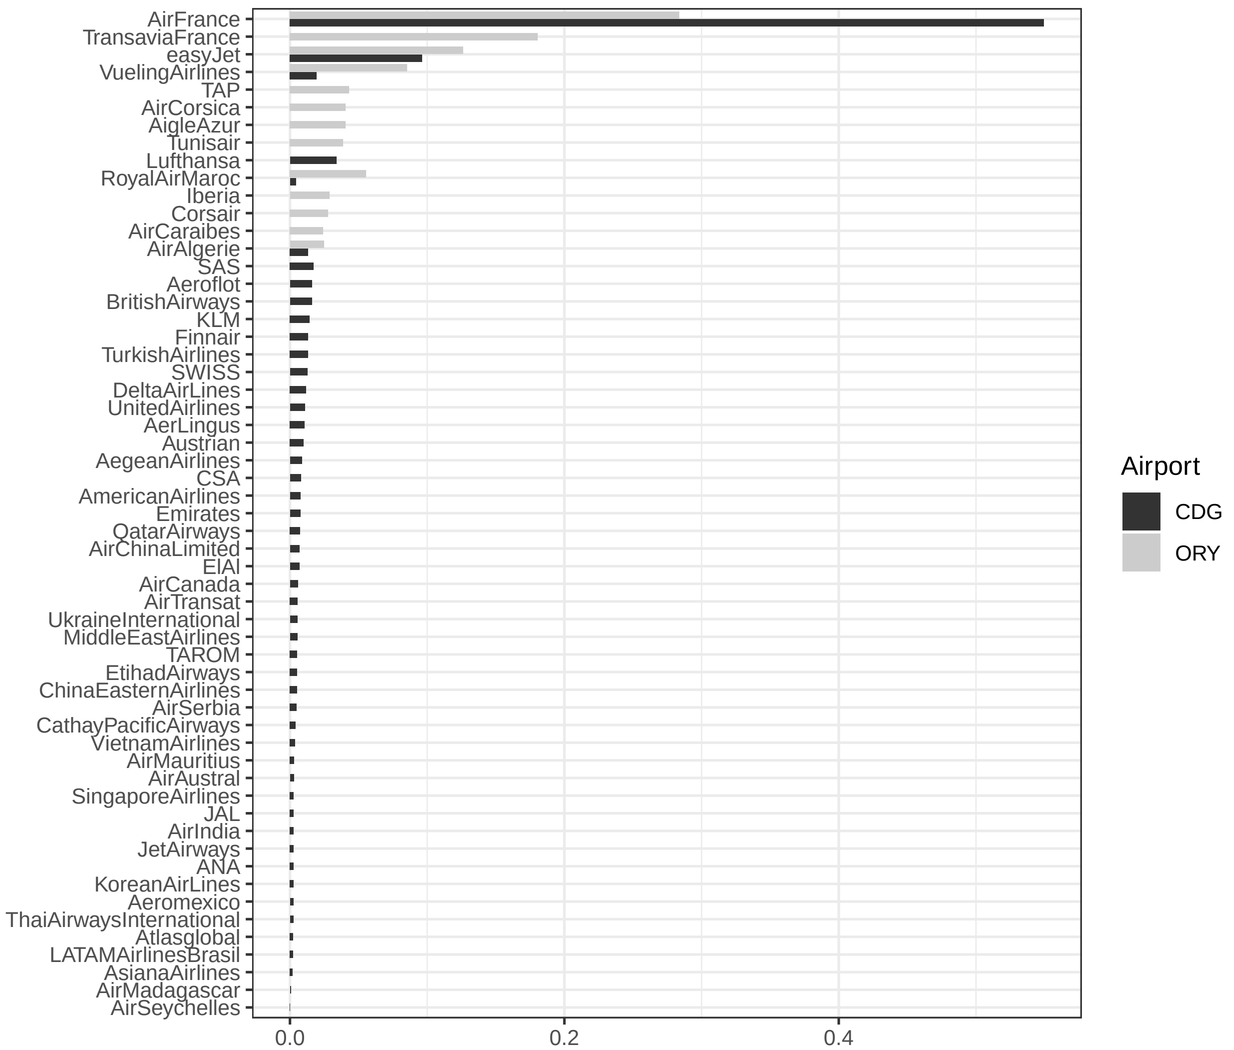

Figure C3: Share of departure slots by airline

Table C7: Routes with connection at CDG

departureAirport	arrivalAirport	Share	ShareCDG	ShareORY	Number	NumberCDG	NumberORY
ISB	BCN	1.00	1.00	0.00	31	31	0
LHE	BCN	1.00	1.00	0.00	31	31	0
MXP	ISB	1.00	1.00	0.00	31	31	0
MXP	LHE	1.00	1.00	0.00	31	31	0
FUE	RNS	1.00	1.00	0.00	3	3	0
NTE	DBV	0.01	0.01	0.00	1	1	0

Notes. Source: Cirium global schedules file, for S18, listing only the operating carrier and no codesharing partners and excluding non-commercial flights. No flights are connecting at ORY.

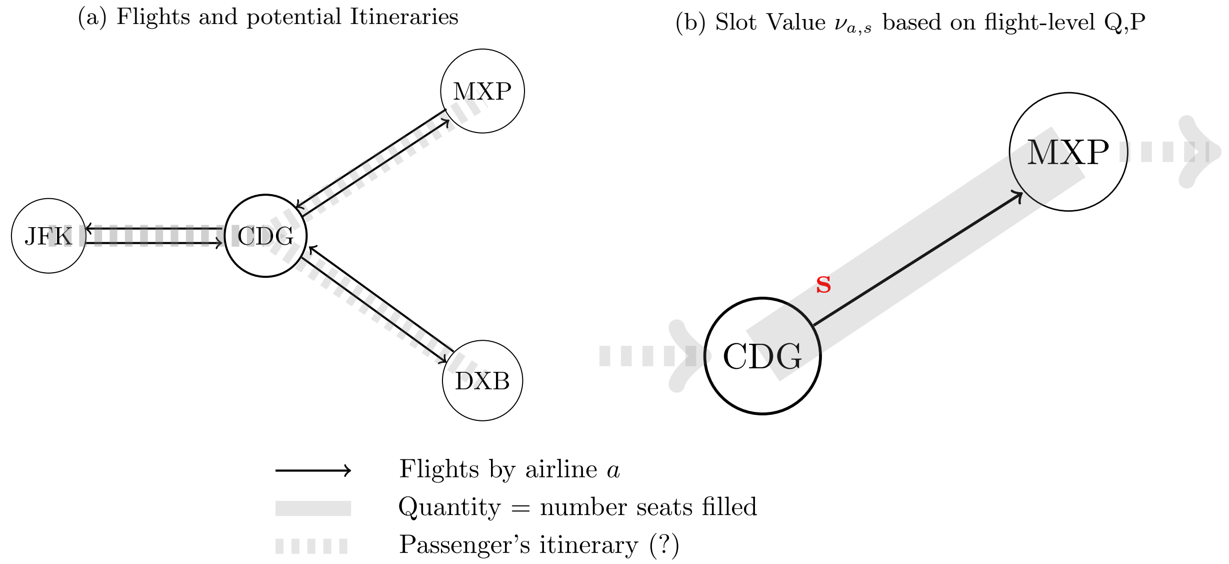

Figure C4: Flight-level analysis to recover the value of slot s

Notes. For example, consider two passengers, one that buys a return ticket to fly from New York (JFK) to Milan Malpensa (MXP) via Paris CDG and another one that buys a return ticket to fly from JFK to Dubai International Airport (DXB) via CDG. They are in the same plane departing from JFK to CDG. From the perspective of a consumer, JFK-CDG-MXP and JFK-CDG-DXB are two different itineraries both with a connecting stop at CDG. The itinerary-level approach that is dominant in the literature would consider these as four directional markets (New York-Milan, Milan-New York, New York-Dubai, Dubai-New York) or two non-directional ones. Demand is modelled as the indirect utility of purchasing a ticket with a dummy for whether the itinerary includes a connection or not. However, from the perspective of an airline carrier, three flights are involved in each direction. The flight that arrives at CDG from JFK is not (necessarily) linked to either of the endpoints of the two itineraries, and the aircraft might fly back to JFK or it might continue to another destination. For this exact reason, the global schedules file has a much lower share of connecting flights at CDG and ORY than the share of departing passengers that are reported to be connecting at these airports.

Table C8: Analysis of demand and supply instruments

	First Stage: OLS of Price on				OLS of Q on
	(1)	(2)	(3)	(4)	(5)
Fuel Consumption	1.292***	0.638***	0.246***	0.368***	
	(0.002)	(0.004)	(0.009)	(0.016)	
Business Class share		1.527***	-1.536***	-0.143	
		(0.095)	(0.105)	(0.212)	
Wet Lease		-0.019	0.079***	-0.138***	
		(0.017)	(0.019)	(0.041)	
Monopoly Markets (at Destination)		1.484***	0.492***	-0.147	59.659***
		(0.041)	(0.049)	(0.108)	(1.027)
Duopoly Markets (at Destination)		0.727***	0.228***	-0.215***	29.098***
		(0.025)	(0.028)	(0.060)	(0.644)
Number of products in Market		-0.001***	0.000***	0.002***	0.013***
		(0.000)	(0.000)	(0.000)	(0.002)
Week of departure		-0.005***	-0.002***	-0.003***	0.050*
		(0.000)	(0.000)	(0.001)	(0.021)
Spacing		0.011***	0.059***	0.041***	
		(0.002)	(0.003)	(0.005)	
Hub (at Destination)		-0.247***	0.361***	0.809***	
		(0.012)	(0.074)	(0.147)	
Nr. Flights (at Destination)		0.005	-0.139***	-0.182***	
		(0.007)	(0.007)	(0.015)	
Number own flights from CDG in Market		0.001***	0.000	-0.003***	-0.117***
		(0.000)	(0.000)	(0.001)	(0.004)
Number competing flights from CDG in Market		0.002***	0.002***	0.003***	-0.025***
		(0.000)	(0.000)	(0.001)	(0.006)
Number competing flights in Market		0.037***	0.036***	0.002	-7.786***
		(0.003)	(0.003)	(0.007)	(0.143)
Nr. own flights in 2 hour window in Market		-0.025**	-0.047***	-0.022	-2.914***
		(0.009)	(0.008)	(0.018)	(0.299)
Nr. competing flights in 2 hour window in Market		-0.019*	-0.008	-0.016	11.510***
		(0.007)	(0.007)	(0.013)	(0.334)
Total capacity own flights in Market		0.001***	0.001***	0.001***	0.155***
		(0.000)	(0.000)	(0.000)	(0.002)
Total capacity competing flights in Market		0.001***	0.000***	0.000***	0.231***
		(0.000)	(0.000)	(0.000)	(0.002)
Num.Obs.	120039	120039	120039	48787	120039
R2	0.785	0.879	0.913	0.890	0.886
R2 Adj.	0.785	0.879	0.912	0.889	0.886
AIC	486992.5	418552.9	379407.9	174727.7	1344878.5
Log.Lik.	-243494.254	-209216.443	-189525.963	-87186.868	-672422.270
Exogenous variables X included		✓	✓	✓	
Destination fixed effects included			✓	✓	
Sample with raw fares only				✓	
Logit IV Model diagnostics:					
– Weak instruments test (p-value)	0	0	0	0	
– Wu-Hausman endogeneity test (p-value)	0	0	0	0	
– Sargan overidentification test (p-value)		0	0	0	

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

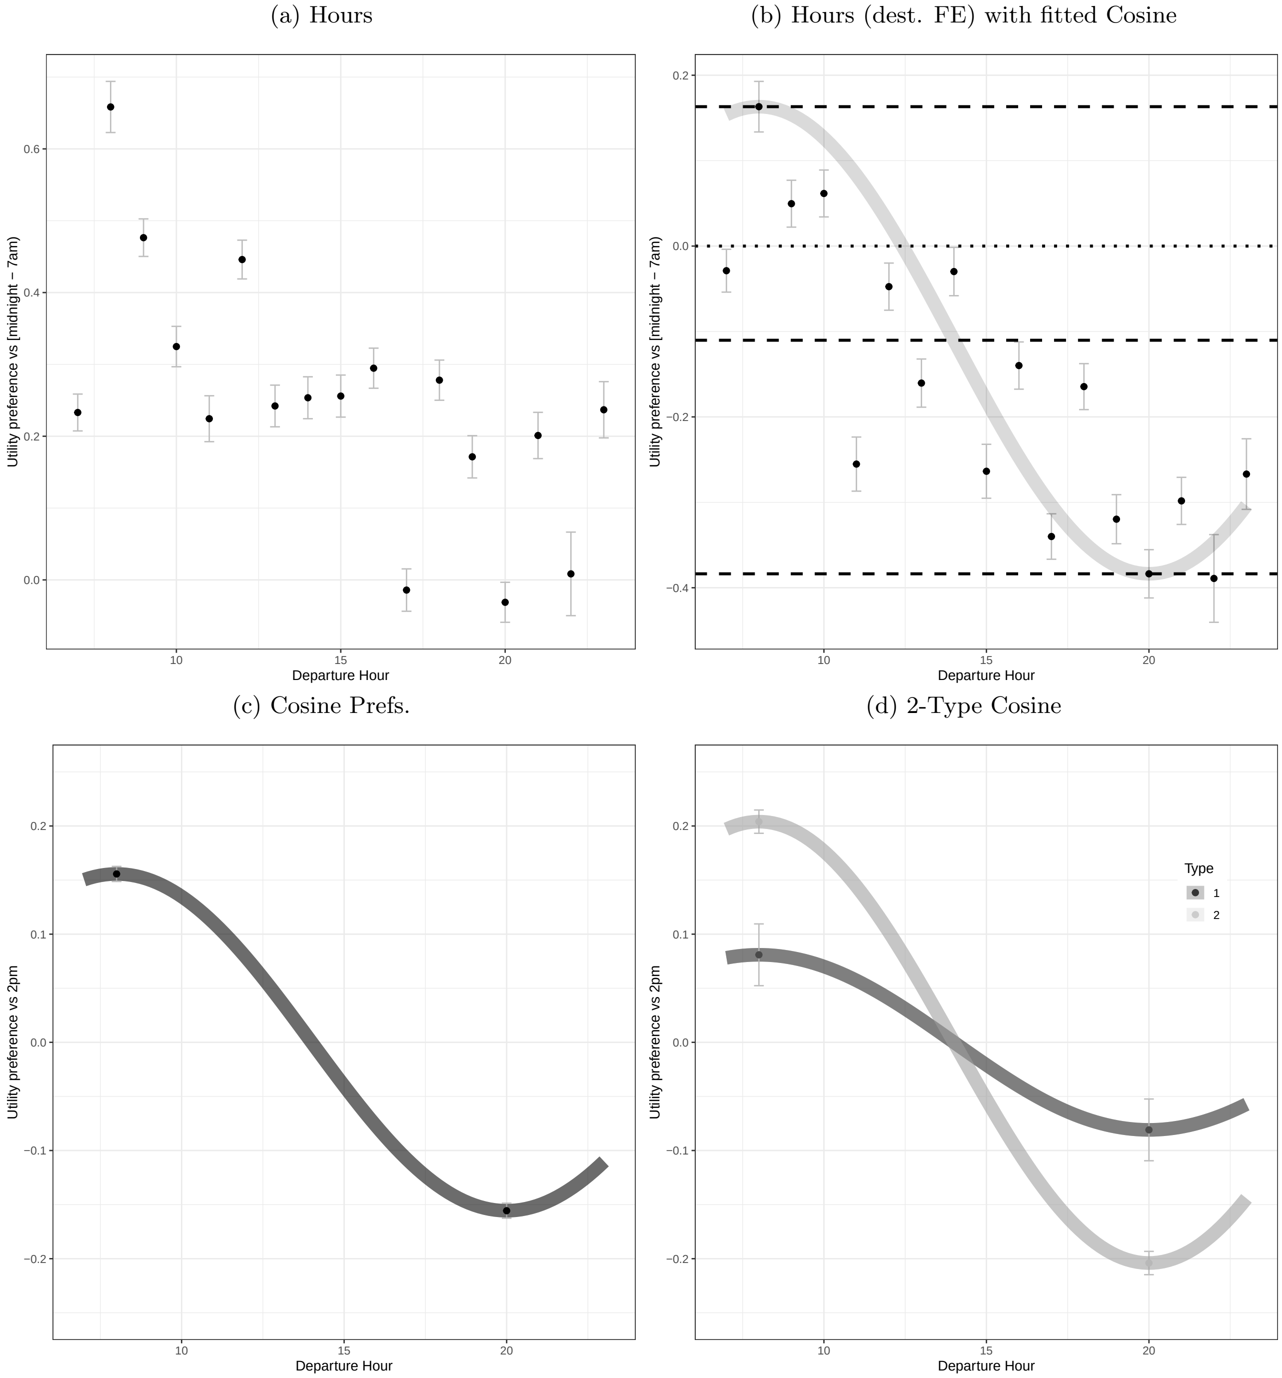

Figure C5: Departure time preferences in different demand specifications

Notes. Plot (a) is based on the demand model with departure hour dummies and without destination fixed effects. It reports the point estimates and 95% confidence interval of the departure hour dummies. The coefficients should be interpreted in relation to the reference group of departing after midnight and before 7am. Departure times are binned into departure hours after midnight — for instance, the departure hour of a flight leaving at 13:46 (1:46 pm) is 13. Plot (b) is based on the departure hour demand model with destination fixed effects. The estimated departure time preferences are also represented by a fitted cosine function $f(x) = a\cos(\frac{2\pi}{24}(x-8))$, with $a\cos(bx)$ the cosine function with amplitude a and period $\frac{2\pi}{b}$ applied to x. As detailed in the text, this function is chosen for its suitability to represent time preferences and due to its fit to the estimated parameters. The amplitude is half the difference between the minimum and maximum of the function, indicated in plot (b) as the distance between the top and middle (or bottom) horizontal dashed lines. Setting $b = \frac{2\pi}{24}$ forces the preferences to repeat every 24 hours. Estimated Cosine-shaped departure time preferences in a model with destination fixed effects are given in plot (c). Plot (d) displays the estimated type-specific Cosine-shaped departure time preferences. The four models are further evaluated in Table 1.

Table C9: Estimated demand and cost parameters

Demand variables		Cost variables	
Nested logit coefficient (kappa)	0.783***	Constant	0.243***
	(0.003)		(0.015)
Population share (type 1)	0.045***	Hub	0.588***
	(0.002)		(0.014)
Fare (type 1)	-0.274***	Spacing	0.087***
	(0.008)		(0.002)
Fare (type 2)	-1.245***	Nr. Flights	-0.143***
	(0.007)		(0.012)
Departure Airport = CDG (type 1)	-1.55***	Hub at Destination	0.085***
	(0.05)		(0.006)
Departure Airport = CDG (type 2)	0.677***	Spacing at Destination	0.021***
	(0.011)		(0.001)
Cosine Departure Time (type 1)	0.081***	Nr. Flights at Destination	0.105***
	(0.015)		(0.003)
Cosine Departure Time (type 2)	0.204***	Fuel Consumption	0.271***
	(0.005)		(0.004)
Day = Tuesday	0.118***	Wet Lease	0.048***
	(0.01)		(0.012)
Day = Wednesday	-0.244***	Departure Airport = CDG	0.343***
	(0.011)		(0.01)
Day = Thursday	0.161***	Slot Constrained Destination	-0.173***
	(0.01)		(0.008)
Day = Friday	-0.117***	Business Class share	1.093***
	(0.01)		(0.065)
Day = Saturday	-0.054***	Flight = Medium Haul	2.127***
	(0.011)		(0.015)
Day = Sunday	0.012	Flight = Long Haul	2.24***
	(0.01)		(0.023)
Air France	-0.136***	Arrival = Mid Day	0.235***
	(0.014)		(0.009)
easyJet	-0.56***	Arrival = End Day	0.033***
	(0.012)		(0.008)
Other Flag Carrier	0.133***	Capacity	-0.001***
	(0.013)		(0)
Lufthansa	-0.388***		
	(0.02)		
Vueling Airlines	-0.147***		
	(0.014)		
Hub	-0.105***		
	(0.011)		
Nr. Flights to Destination	0.24***		
	(0.032)		
Nr. Flights to Destination (squared)	-0.174***		
	(0.023)		
Connectivity Arrival Airport	0.211***		
	(0.01)		
Connectivity Departure Airport	0.609***		
	(0.011)		

Notes. The estimated destination fixed effects are summarized in Table C10. The *Fare* is the average ticket price in €100. Just like prices, costs are also reported in €100. Summary statistics are given in Table C11 in the online appendix.

Table C10: Estimated destination fixed effects

City	Airport	FE
Highest utility (vs outside options, incl. train)		
Libreville	LBV	-0.12
Atlanta	ATL	-0.04
Abidjan	ABJ	0.00
Cayenne	CAY	0.26
Philipsburg	SXM	0.32
Higuey	PUJ	0.69
Mahe Island	SEZ	0.80
Port Louis	MRU	0.92
Saint-Denis	RUN	1.73
Dzaoudzi	DZA	2.73
Lowest utility (vs outside options, incl. train)		
London	LTN	-7.20
London	LHR	-6.88
Hamburg	HAM	-6.08
Liverpool	LPL	-6.04
Milan	LIN	-5.96
Berlin	TXL	-5.76
Zurich	ZRH	-5.76
Milan	MXP	-5.76
Marseille	MRS	-5.75
Barcelona	BCN	-5.72

Notes. The demand model includes fixed effects for all 119 destinations (airports) in the data. The 10 highest and lowest destination fixed effect estimates are reported, reflecting destinations where the utility of flying is estimated to be the most and least attractive relative to the outside option, respectively.

Table C11: Summary statistics of selected demand and cost variables

	Mean	Min	Median	Max
Fare	3.02	0.39	2.07	19.27
Quantity	175.85	4.00	163.00	615.00
CDG	0.67	0.00	1.00	1.00
Hub	0.85	0.01	0.40	1.89
FlightstoDest	0.40	0.10	0.30	1.50
Spacing	1.42	0.07	0.41	13.75
Flights	0.91	0.01	0.38	2.36
HubDest	1.30	0.08	1.20	3.10
SpacingDest	2.95	0.00	1.03	16.83
FlightsDest	0.73	0.00	0.13	19.44
FuelBurnperSeatL9	1.81	0.23	0.91	14.78
WetLease	0.08	0.00	0.00	1.00
SlotConstrained	0.65	0.00	1.00	1.00
FirstBusiness	0.04	0.00	0.00	0.42
MediumHaul	0.08	0.00	0.00	1.00
LongHaul	0.14	0.00	0.00	1.00
ArrivalMidDay	0.33	0.00	0.00	1.00
ArrivalEndDay	0.48	0.00	0.00	1.00
Capacity	205.75	48.00	180.00	615.00
DepHour	13.76	5.00	14.00	23.00
Total number of products (flights)	120039.00			
Total number of flights by Air France	63793.00			
Total number of markets (destination-months)	952.00			
Average number of products per market	279.27			
Total number of firms	57.00			
Average number of firms per market	3.27			

Table C12: Changes in slot ownership - largest firms participate

		Post-auction				
		AirFrance	easyJet	Lufthansa	VuelingAirlines	TOTAL pre-auction
Pre-auction	AegeanAirlines	2	0	0	0	2
Pre-auction	AerLingus	4	0	0	0	4
Pre-auction	Aeroflot	4	0	0	0	4
Pre-auction	AigleAzur	1	0	0	0	1
Pre-auction	AirAlgerie	2	1	0	1	4
Pre-auction	AirAustral	1	0	0	0	1
Pre-auction	AirCanada	2	0	0	0	2
Pre-auction	AirCaraibes	2	2	0	0	4
Pre-auction	AirCorsica	1	3	0	0	4
Pre-auction	AirFrance	72	5	4	4	85
Pre-auction	AirIndia	1	0	0	0	1
Pre-auction	AirTransat	1	0	0	0	1
Pre-auction	ANA	1	0	0	0	1
Pre-auction	Corsair	1	1	0	0	2
Pre-auction	DeltaAirLines	1	0	0	0	1
Pre-auction	easyJet	14	4	3	2	23
Pre-auction	ElAl	1	0	0	0	1
Pre-auction	Emirates	1	0	0	0	1
Pre-auction	EtihadAirways	0	1	0	0	1
Pre-auction	Finnair	1	0	1	0	2
Pre-auction	Iberia	1	1	0	0	2
Pre-auction	LATAMAirlinesBrasil	0	1	0	0	1
Pre-auction	Lufthansa	6	0	1	0	7
Pre-auction	MiddleEastAirlines	2	0	0	0	2
Pre-auction	QatarAirways	1	1	0	0	2
Pre-auction	RoyalAirMaroc	6	2	0	0	8
Pre-auction	SAS	1	0	0	0	1
Pre-auction	SWISS	3	0	0	0	3
Pre-auction	TAP	4	1	0	0	5
Pre-auction	TransaviaFrance	9	7	0	2	18
Pre-auction	Tunisair	3	1	0	0	4
Pre-auction	TurkishAirlines	0	0	0	1	1
Pre-auction	UkraineInternational	4	0	2	1	7
Pre-auction	UnitedAirlines	1	0	0	0	1
Pre-auction	VietnamAirlines	2	0	0	0	2
Pre-auction	VuelingAirlines	4	2	0	1	7
NA	TOTAL post-auction	160	33	11	12	216

Table C13: Changes in slot ownership - all firms participate

		Post-auction													
		AirAustral	AirCanada	AirCaraibes	AirFrance	AirMauritius	Corsair	Emirates	Iberia	QatarAirways	RoyalAirMaroc	TAP	TransaviaFrance	Tunisair	TOTAL pre-auction
Pre-auction	AegeanAirlines	0	0	0	2	0	0	0	0	0	0	0	0	0	2
Pre-auction	AerLingus	0	0	0	4	0	0	0	0	0	0	0	0	0	4
Pre-auction	Aeroflot	0	0	0	5	0	0	0	0	0	0	0	0	0	5
Pre-auction	AigleAzur	0	0	1	0	0	0	0	0	0	0	0	0	0	1
Pre-auction	AirAlgerie	0	0	0	4	0	0	0	0	0	0	0	0	0	4
Pre-auction	AirAustral	0	0	0	1	0	0	0	0	0	0	0	0	0	1
Pre-auction	AirCanada	0	0	0	2	0	0	0	0	0	0	0	0	0	2
Pre-auction	AirCaraibes	0	0	2	1	0	0	0	1	0	0	0	0	0	4
Pre-auction	AirCorsica	0	0	2	1	0	0	0	0	0	0	0	0	1	4
Pre-auction	AirFrance	1	1	3	54	0	2	7	7	0	2	1	17	2	97
Pre-auction	AirIndia	0	0	0	1	0	0	0	0	0	0	0	0	0	1
Pre-auction	AirTransat	0	0	0	1	0	0	0	0	0	0	0	0	0	1
Pre-auction	ANA	0	0	0	1	0	0	0	0	0	0	0	0	0	1
Pre-auction	Corsair	0	0	0	1	0	0	0	1	0	0	0	0	0	2
Pre-auction	DeltaAirLines	0	0	0	1	0	0	0	0	0	0	0	0	0	1
Pre-auction	easyJet	0	0	3	13	1	0	3	4	0	1	1	0	0	26
Pre-auction	ElAl	0	0	0	0	0	0	1	0	0	0	0	0	0	1
Pre-auction	Emirates	0	0	0	1	0	0	0	0	0	0	0	0	0	1
Pre-auction	EtihadAirways	0	0	0	1	0	0	0	0	0	0	0	0	0	1
Pre-auction	Finnair	0	0	0	2	0	0	0	0	0	0	0	0	0	2
Pre-auction	Iberia	0	0	1	1	0	0	0	0	0	0	0	0	0	2
Pre-auction	LATAMAirlinesBrasil	0	0	0	1	0	0	0	0	0	0	0	0	0	1
Pre-auction	Lufthansa	0	0	0	7	0	0	0	0	0	0	0	0	0	7
Pre-auction	MiddleEastAirlines	0	0	0	1	0	0	1	0	0	0	0	0	0	2
Pre-auction	QatarAirways	0	0	0	3	0	0	0	0	0	0	0	0	0	3
Pre-auction	RoyalAirMaroc	0	0	2	5	0	0	0	2	0	0	0	1	0	10
Pre-auction	SAS	0	0	0	0	0	0	1	0	0	0	0	0	0	1
Pre-auction	SWISS	0	0	0	3	0	0	0	0	0	0	0	0	0	3
Pre-auction	TAP	0	0	2	4	0	0	0	1	0	0	0	0	0	7
Pre-auction	TransaviaFrance	0	0	7	4	0	0	0	5	0	2	0	2	0	20
Pre-auction	Tunisair	0	0	1	1	0	0	0	2	0	0	0	0	0	4
Pre-auction	TurkishAirlines	0	0	0	0	0	0	1	0	0	0	0	0	0	1
Pre-auction	UkraineInternational	0	0	0	7	0	0	0	0	0	0	0	0	0	7
Pre-auction	UnitedAirlines	0	0	0	0	0	0	0	0	1	0	0	0	0	1
Pre-auction	VietnamAirlines	0	0	0	2	0	0	0	0	0	0	0	0	0	2
Pre-auction	VuelingAirlines	0	0	0	4	0	0	0	0	0	0	0	5	1	10
NA	TOTAL post-auction	1	1	24	139	1	2	14	23	1	5	2	25	4	242