A Market for Airport Slots

Marleen Marra*

March 26, 2026

Abstract

Globally, scarce airport capacity is rationed by awarding time-specific slots free of charge to airlines, incentivising retention regardless of efficiency. I quantify the efficiency losses and welfare effects of market-based reallocation, defining slot values as equilibrium profits from a structural flight-level model. Reallocating 6% of weekly slot pairs in two airports generates €213M in consumer surplus through substitution of under-utilised slots for long-haul flights at suitable hours; fare effects are negligible. Larger operations yield cost and demand advantages and reduce outside hubs as substitutes. This reinforces concentration, but unlike in standard markets, the resulting network convenience benefits consumers at the margin.

JEL: L13, L93, D44, C57

Keywords: *Airline industry, Structural analysis, Airport slots, Auctions, Network effects, Spillovers, Firm heterogeneity, Time preferences*

KU Leuven, Naamsestraat 69, 3000 Leuven, Belgium, e-mail: marleen.marra@kuleuven.be. I would like to thank Peter Cramton, Adam Dearing, and Kevin Williams for insightful discussions of the paper, and 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 Sam Altmann, Nicolas Astier, Christian Bontemps, Moshe Buchinsky, Clément de Chaisemartin, Allan Crawford, Pierre Dubois, Silke Forbes, Jeremy Fox, Phil Haile, 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 to other seminar participants at MIT, Harvard Univeristy, Columbia University, University of Berkeley, HEC, INSEAD, Paris School of Economics, Toulouse School of Economics, Tilburg University, Crest, NHH Bergen, Mannheim University, Queen Mary University, Carlos III, University of Geneva, KU Leuven, Cergy Paris, and Sciences Po, and participants at the 2026 SoAR Symposium, the 2025 Alpine IO Conference, Cowles Conference 2024, Rice Structural Econometrics 2024, AFSE 2023, Leuven Summer Event 2022, MACCI Summer Institute 2022, CMES 2021, and EARIE 2020 for feedback and useful conversations. 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 from 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.

Introduction

The globally adopted system for allocating scarce airport capacity awards time-specific take-off and landing rights ("airport slots") free of charge to incumbent carriers based on historic use, covering over 200 airports and about 50% of passengers worldwide.[1] Airlines retain their departure rights from the previous year as long as they cancel less than 20% of associated flights. The airport slot policy has garnered considerable attention because of its implications for airline competition, market efficiency, and potential windfall profits for slot-owning airlines. In discussions on noise reduction, pollution control, and sustainable aviation, the need for a more efficient allocation policy is becoming increasingly apparent.

Transitioning to auctions may not enhance welfare. First, a firm's valuation of a slot depends on the profits it can generate by operating flights in that slot, which need not align with consumer value. Second, since slots are not tied to specific routes or flights, different airlines can use them for different markets, disrupting the alignment between allocative and productive efficiency. Third, consumers derive benefits from carriers holding a larger share of slots through associated network perks — shorter layovers, better connections, higher frequency — as widely documented in the literature.[2] Whether such network benefits outweigh reduced competition is uncertain ex ante, and the direction is not obvious without simultaneously quantifying the product-variety and network-externality channels. This paper does so and finds that reallocation raises welfare.

While studying this issue, the paper makes three contributions. First, it develops and estimates a flight-level structural model using novel data on global flights from France, identifying consumer preferences over departure and arrival times and flight-level network spillovers in costs and demand. These features are essential to capture why some slots command substantially higher values than others — and why not all slots are contested even at the most congested level-3 airports. Time preferences cause carriers to earn higher equilibrium mark-ups on flights at favourable hours; network spillovers determine how much a slot is worth to a carrier's wider network. Without them, it would be hard to rationalise prices such as the $75 million paid by Oman Air in 2016 for a single Heathrow slot pair — an early morning arrival connecting businessmen onward to the US and Europe — or to predict which slots are contested and which are not. The estimated time preferences imply that business travellers pay €145 more for a morning arrival relative to early afternoon.

The model is estimated on prices and quantities of all direct flights from Paris Charles de Gaulle (CDG) and Paris Orly (ORY) to 119 global destinations — novel data for this type of analysis, as the literature predominantly relies on U.S. route-level aggregates. Unlike

[1] Source: IATA. Section I describes the allocation policy in detail.
[2] See, e.g., Borenstein (1989), Berry (1990), Brueckner, Dyer and Spiller (1992), Berry and Jia (2010).

the typical itinerary-level demand model, demand is specified at the flight level, where consumers value departure timing and flights are linked through a connection technology. Connecting passengers — 21.7% of CDG departure traffic — are unobserved as itineraries but captured through connectivity premia at origin and destination airports. Business and leisure travellers are willing to pay €77 and €24, respectively, for one standard deviation more connecting flights at the departure airport — roughly three times their valuation of equivalent connectivity at the destination, reflecting Paris' hub function for the dominant carrier in the sample. On the supply side, the model recovers flight-level marginal costs accounting for scheduling efficiencies and network economies.

The second contribution develops a new approach to valuing slots without bid data. Slot values are defined as the incremental profit a carrier obtains by operating a flight in the slot relative to the counterfactual where its worst-case competitor wins instead, capturing the carrier's full equilibrium profits including network spillovers. Assuming carriers use additional slots for their best-performing feasible flights, and that this choice is public information, makes the slot value problem tractable and yields a private values structure (Jehiel and Moldovanu, 2001), guaranteeing the existence of an efficient incentive-compatible mechanism despite allocative externalities. Analysing flights in isolation would overstate inefficiency and ignore the industry's characterizing hub strategies. For example, Air France's marginal flight earns profits eight times below the sample average, which is consistent with use-it-or-lose-it retention of slots regardless of direct-route value,[3] but misses that marginal flights might support profitable long-haul connections. To address this, direct spillovers are recovered from the structural model and hub opportunity costs are proxied by the within-carrier-hour gap between average and minimum estimated flight profit, validated by its strong correlation with carriers' broader flight networks and hub dependence outside Paris.

Estimated slot values are consistent in magnitude with observed Heathrow transaction prices. Three-quarters of the variation across all carrier–slot combinations is across carriers within the same slot, reflecting different network opportunities rather than slot timing — an additional 20% varies across hours of the day. Air France commands higher values for many slots, attributable to network-related cost and demand advantages and larger expansion opportunities: connectivity premia, scheduling efficiencies, and a wider set of feasible flight additions. The median spread between the highest and second-highest value is 133%, lending credibility to the welfare estimates — unmodelled factors would need to be both large and systematically correlated with carrier identity to yield very different outcomes.

The third contribution presents the first welfare assessment of market-based slot allocation linking slot values to flight-level profitability under equilibrium conditions. Reallocating

[3] *Slot hoarding*, see, e.g., De Wit and Burghouwt (2008), Fukui (2012), Kösters, Meier and Sieg (2023).

245 weekly slot pairs — approximately 6% of all pairs in the sample season — generates €213M in consumer surplus and €269M in variable profits. Gains are driven by short-to-long-haul repurposing: 55% of slot changes redirect a short-haul flight to a long-haul route, and long-haul markets account for €165M of the consumer surplus gain. Morning slots generate 49% of gains despite representing 29% of auctioned slots. Fare increases are negligible (€0.46 per flight on average), confirming that gains come from improved product offerings rather than price effects, and consistent with the incremental nature of the reallocation leaving market structures largely intact. A restricted specification imposes that network changes correspond to observed changes in Paris between the summer of 2018 and 2019, reflecting additional constraints like unmodeled destination slot scarcity, return-flight requirements, and fleet limitations. The welfare effects remain positive, but limiting expansion opportunities results reduces the consumer surplus gain to €121M.

Whether firms are overall better off depends on the redistribution of auction rents, but as aggregate welfare gains are positive, a well-designed mechanism can make all participants at least indifferent. These findings speak to the central ambiguity (Borenstein, 1988): efficient licence markets do not necessarily improve welfare, but empirically they do here.

This research connects two pillars of empirical Industrial Organization: demand-supply models for posted-price markets (e.g., Berry, Levinsohn and Pakes, 1995; Berry and Haile, 2014) and bidding models for auctions (e.g., Guerre, Perrigne and Vuong, 2000). The auction literature identifies bidders' latent values from observed bids. This paper instead defines value directly as profit-generating potential, estimated via demand and supply — an approach essential for analysing the welfare effects of *introducing* a market mechanism when no revealed-preference data exist. To my knowledge, Doraszelski et al. (2017) is the only prior ex-ante analysis of this kind, though not basing values on equilibrium estimates.

While Levine (1986) already identified airport slots as important entry barriers, most empirical IO abstracts from them, partly due to data-driven focus on largely unconstrained US markets. Exceptions include Ciliberto and Williams (2010) on gate restrictions and Park (2020) on slot divestitures. Other work examines slot policies without modelling industry equilibrium responses (e.g., Borenstein, 1989; Odoni, 2020), or addresses computational feasibility of market mechanisms (e.g., Rassenti, Smith and Bulfin, 1982). This paper is the first to quantify welfare effects of market-based slot allocation under equilibrium conditions. Bauer (2025) extends this framework to a broader sample of European airports using route-carrier aggregates and a genetic algorithm for slot assignment, contributing to a growing body of research on the European airline industry (e.g., Bergantino et al., 2024; Bontemps, Martini and Porta, 2026; Chen, Chen and Nesheim, 2026). Adler, Andreana and De Jong (2026) study the societal value of airport slots within an operations research framework.

Lazarev (2013), Aryal, Murry and Williams (2024), and Williams (2022) are closely related studies that estimate flight-level structural models to move beyond route-level analysis — in their case, to study dynamic pricing. The disaggregated model here is designed to capture the relative attractiveness of different slots: consumer preferences for flight timing, network position, and flight-level costs. Unlike the typical itinerary-level demand model, demand is specified at the flight level, where consumers value departure timing and flights are linked through a connection technology adapted from Yuan and Barwick (2024). This formulation is essential for evaluating slot-specific policies, as it allows network modifications to generate spillovers across products and markets, and captures that timing and network position are first-order drivers of slot value differences.

Finally, the paper contributes to the literature on auctions with aftermarkets, where welfare effects of efficient allocation are theoretically uncertain when winners compete downstream (e.g., Janssen and Karamychev, 2009; Ershov and Salant, 2022). The existing literature focuses on spectrum auctions, where licences reduce marginal costs across markets. Airport slots differ in two respects: network effects operate on both demand and supply sides, and the slot does not specify which product it applies to — carriers choose which flight to operate (cf. Borenstein, 1988). This connects to work on endogenous product choice preceding price competition (e.g., Wollmann, 2018; Fan and Yang, 2020).

The paper proceeds as follows. Section I describes the slot allocation process and data, and presents model-free evidence of inefficiency. Section II presents the flight-level demand and supply model, identification, and estimation. Section III discusses estimation results. Section IV links flight profits to slot values and details counterfactual simulations. Section V presents welfare impacts of market-based reallocation. Section VI concludes.

I Industry and data description

As of 2024, half of all airline passengers worldwide depart from slot-constrained airports, where demand for runway and gate access exceeds capacity.[4] At these airports, airlines must obtain slots from a coordinator to operate. The allocation process, governed by IATA's World Airport Slot Guidelines (and in Europe by Council Regulation EEC 95/93), has two key features. First, incumbents retain slots from the previous season if utilized at least 80% of the time (the "use-it-or-lose-it" rule). Second, slots that are returned, newly created, or forfeited enter a pool, from which coordinators fulfill waitlist requests—with 50% reserved for new entrants, typically carriers owning at most six (arrival+departure) slots that day.

This administrative system allocates slots based on historical precedent and queuing priority rather than willingness to pay. In Europe, the initial allocation following market

[4]IATA statistics: `https://www.iata.org/en/programs/ops-infra/slots/`.

liberalization in 1997 favored previously state-owned flag carriers, who retain dominant positions in major airports today. Secondary trading is permitted in some jurisdictions but remains thin. This study aims to quantify the welfare effects of moving toward a market-based allocation, based on how much carriers value the slots. The focus of the study is two airports: Paris Charles-de-Gaulle (CDG) and Paris Orly (ORY). Both airports are among the most slot-constrained international airports in the world, and CDG is the second-busiest airport in Europe after Heathrow Airport (Milioti and Odoni (2024)), serving more than 30 million passengers per year in 2019. At Paris Orly, the 2019 Aigle Azur bankruptcy released nearly 10,000 slots midseason; yet only 18% of new entrant requests could be accommodated, illustrating the magnitude of excess demand. In Appendix A and Marra (2024) I provide further institutional details.

I.A Data sources

Two main data sources are used. The first is data from the French civil aviation regulator DGAC, containing a sample of ticket prices and seats occupied for return flights departing from CDG or ORY in the 2018 summer season (S18, between 26 March and November 3rd). The administration uses these data to construct the travel component of the consumer price index. The unique advantage of this data is that it contains flight-level identifiers including departure time and date—information unavailable in the US Department of Transportation's DB1B sample typically used in airline research or in OAG's quarterly flight aggregates.

The price variable is the average lowest offered fare, weighted by passenger surveys at CDG and ORY to reflect actual purchase prices. Table 1 shows that DGAC quantities compare well with DB1B data: destinations appearing in both datasets have similar passenger traffic shares. However, the DGAC sampling algorithm does not condition on departure time, resulting in missing fares for flights when a lower-priced alternative departed at a different time that day.[5]

Since prices reflect average return fares for direct flights, half is attributed as the marginal revenue contribution per passenger. This rules out that the marginal revenue contribution from passengers connecting in CDG is different from the direct fare of the flight from CDG. This proportionality principle holds in a sample of U.S. ticket data (Williams, 2022) as well as in the sample of US-CDG flights in the SIAT survey data used in Aryal, Murry and Williams (2024), and is routinely applied to the raw DB1B data.[6]

[5] I impute missing fares using random forest, achieving $R^2 = 0.89$ in the held-out test sample; details on the imputation procedure and its consistency properties are provided in Marra (2024). When the flight-level unobservable ξ is constant across offers and independent of offer-specific heterogeneity, data are missing at random conditional on observables and ξ, and imputation yields consistent estimates (Dagenais, 1973).

[6] The mean fares accounted to the direct flights in Aryal, Murry and Williams (2024) are not statistically different at the 5% level when the itinerary includes a connection or not. I thank the authors for sharing this information.

Table 1: Quantity in DGAC and T100 data: comparing passenger shares by destination

	ATL	BOS	EWR	IAD	JFK	LAX	MIA	ORD	SFO	all
DGAC	0.12	0.07	0.06	0.11	0.27	0.11	0.07	0.08	0.12	1.00
T100	0.11	0.07	0.10	0.08	0.27	0.14	0.06	0.06	0.10	1.00

The second data source, purchased from Cirium, contains the universe of scheduled flights globally for January 2018–October 2019. These data are not limited to Paris-destination pairs and enable analysis of how flight networks affect slot values. I also compare observed network changes between the summer season of 2019 (S19) and S18 to validate my model estimates and to comment on the role of system-wide slot constraints.

Finally, weekly jet fuel prices, city populations, and airport coordinates are obtained from the US Energy Information Agency, Simplemaps, and Google Maps, respectively. Aircraft-specific fuel consumption parameters are taken from Seymour et al. (2020)

I.B Estimation sample and summary statistics

Quantity data are available for all 300 destinations reachable from CDG and ORY by commercial flight, but the DGAC collects prices only for a subset. The 119 destinations in the estimation sample are selected for their importance in terms of carrier count and flight frequency. The most popular destinations for ticket sales are Nice, Toulouse, Barcelona, Madrid, Montreal, Lisbon, and New York. All cities are retained regardless of population size, as smaller cities (such as Toulouse) can still be important in terms of demand volume. The sample is restricted to the 55carriers for which the DGAC samples fares, which are more relevant in the sense of serving more destinations and offering more flights. Ghost flights are not important in the data.[7] Only approximately 0.2% of flights depart with fewer than 25 passengers, and approximately 0.6% with fewer than 25% of seats occupied. Flights with average fares exceeding the 99th percentile and the 49 flights with fewer than 10 seats occupied are excluded. The final estimation sample contains 119,810 flights (observations); summary statistics are presented in Table 3.

The average flight carries 176 passengers, with an average load factor of 0.85; only 5% of flights are full. The average fare is €301, with a large standard deviation reflecting the mix of low-cost short-distance and more expensive long-haul flights. Fare variation within and across departure times is substantial and more pronounced at CDG than at ORY. ORY is a smaller airport located closer to the city and offers more short-haul flights; average flight distance is 1,629 km (1,012 miles) at CDG and 972 km (604 miles) at ORY.

Among the 55 carriers, the Air France Group (Air France, KLM, and Transavia, hence-

[7] A ghost flight is a flight that takes off with no or very few passengers, usually motivated by the 80% "use-it-or-lose-it" requirement of the current slot allocation process. Flights may also take off (near-)empty for other reasons, for instance, to support a hub function at the destination.

Table 2: Networks by carrier

	AF	LH	U2	VY
Panel A: Paris Operations				
Destinations served from Paris	105	2	30	13
Mean flights to destination (J_{atF})	94.33	184.06	61.11	77.21
Panel B: Outside Network Presence				
Mean # flights outside airports to dest. (J_{atO})	540.63	13,512.38	633.44	865.90
% destinations with outside service	71.59	100.00	100.00	100.00
Mean # outside hubs serving dest. (H_{atO})	1.25	7.94	8.84	1.68
% destinations served by outside hub	42.07	100.00	98.74	98.13
Mean capacity on outside flights (Cap_{atO})	123.67	167.56	168.07	178.66

Notes. Outside statistics based on Global Schedule data, for all markets in the estimation sample. The carrier codes are AF=Air France, LH=Lufthansa, U2=easyJet, VY=Vueling Airlines. A hub serves at least 30 destinations by direct flight.

forth "Air France") is the dominant carrier, holding more than half of the slots in the sample at both airports. Air France and Lufthansa both operate hub-and-spoke networks: Lufthansa's hubs are at destinations (Frankfurt and Munich), whereas Air France's primary hub is at the origin (Paris), with secondary hubs at select destinations such as Amsterdam. The next two largest carriers, easyJet and Vueling Airlines, follow a more decentralized strategy, though with disproportionately many departures from the UK and Spain, respectively. Table 2 summarizes network characteristics for these four carriers. Air France dominates Paris operations, serving 105 markets from CDG/ORY, but maintains limited outside network presence with only 1.25 outside hubs per market on average (a hub is an airport serving at least 30 destinations by direct flight). Lufthansa presents the opposite pattern: minimal Paris presence (2 markets) but an extensive outside network with flights from 8 hubs and over 13,5k outside flights per destination. easyJet and Vueling occupy intermediate positions, serving 30 and 13 Paris markets respectively, with considerable outside presence.

I.C Suggestive evidence for inefficiency

The raw data provides initial model-free evidence consistent with inefficient slot allocation and its competitive consequences. On short-haul routes under 400km, Air France operates 63% more flights than competitors yet achieves load factors 5.5 percentage points lower (81.6% vs 87.1%).[8] This is consistent with slot hoarding: under the "use-it-or-lose-it" rule, incumbents maintain high frequency even when demand does not justify it, retaining valuable slots and deterring entry (De Wit and Burghouwt (2008), Fukui (2012), Kösters, Meier and

[8]The following destinations in the data are less than 400km from Paris: Amsterdam, Bordeaux, Bristol, Frankfurt, Geneva, Liverpool, London, Marseille, Milan, Toulouse, and Zurich, with flights offered by: Air France, British Airways, Lufthansa, SWISS, Vueling Airlines, and easyJet.

Table 3: Summary statistics of selected demand and cost variables

	All		Air France		Competitors		
	Mean	SD	Mean	SD	Mean	SD	p
Fare (100 €)	3.01	2.58	3.08	2.61	2.94	2.55	0.111
Quantity	175.99	81.77	172.73	84.10	179.69	78.87	0.008
Network & Connectivity							
Departure = Morning	0.22	0.42	0.26	0.44	0.18	0.38	<0.001
Departure = Mid Day	0.40	0.49	0.38	0.48	0.42	0.49	<0.001
Departure = End Day	0.38	0.48	0.36	0.48	0.40	0.49	<0.001
Nr. Flights to Dest. on Date	0.41	0.30	0.49	0.34	0.32	0.20	<0.001
Connectivity Dep. Airport	0.86	1.04	1.51	1.03	0.13	0.29	<0.001
Connectivity Arr. Airport	0.57	1.01	0.27	0.66	0.92	1.21	<0.001
Spacing	1.34	2.46	0.11	0.05	2.74	3.04	<0.001
Spacing (at Dest.)	3.34	5.75	5.14	6.50	1.28	3.83	<0.001
Nr. Destinations	0.85	0.82	1.45	0.68	0.17	0.21	<0.001
Nr. Destinations (at Dest.)	1.37	0.71	1.34	0.68	1.40	0.73	0.039
Nr. Flights	1.02	0.97	1.81	0.66	0.12	0.12	<0.001
Nr. Flights (at Dest.)	0.73	1.34	0.15	0.50	1.38	1.66	<0.001
Cost Drivers							
Fuel Consumption (100 €)	1.81	2.04	1.73	2.02	1.89	2.05	0.008
Medium Haul	0.08	0.27	0.06	0.24	0.09	0.29	<0.001
Long Haul	0.14	0.35	0.14	0.35	0.15	0.36	0.347
Arrival = Mid Day	0.33	0.47	0.32	0.47	0.34	0.47	0.122
Arrival = End Day	0.48	0.50	0.48	0.50	0.48	0.50	0.883
Capacity	205.79	83.91	200.43	85.51	211.87	81.62	<0.001
Wet Lease	0.08	0.27	0.13	0.34	0.02	0.14	<0.001
Slot Constrained Dest.	0.65	0.48	0.61	0.49	0.70	0.46	<0.001
First or Business class (% of seats)	4.00	6.44	3.07	5.09	5.05	7.56	<0.001
Slot Utilization							
Load Factor (%): <400km	83.44	15.66	81.59	16.55	87.09	12.97	<0.001
Load Factor (%): Other flights	85.59	15.56	86.91	13.66	84.29	17.14	<0.001
Frequency: <400km	0.72	0.33	0.83	0.32	0.51	0.23	<0.001
Frequency: Other flights	0.24	0.14	0.22	0.11	0.26	0.16	<0.001

Market Concentration	All		High AF Share		Low AF Share		
	Mean	SD	Mean	SD	Mean	SD	p
AF Slot Share	0.52	0.30	0.76	0.20	0.28	0.18	<0.001
HHI (Flights)	0.64	0.24	0.70	0.23	0.58	0.24	<0.001
Nr. Competitors	2.04	0.84	1.84	0.74	2.25	0.89	<0.001
Fare (All)	4.64	3.00	5.22	3.12	4.04	2.75	<0.001
Fare (Air France)	4.93	3.02	5.31	3.16	4.42	2.74	<0.001
Fare (Competitors)	3.92	2.73	3.86	2.58	3.96	2.82	0.583

Sample Characteristics	All	CDG	ORY
Nr. products (flights)	119,810	80,879	38,931
Nr. Air France flights	63,705	45,615	18,090
Nr. markets (dest.-months)	952	888	431
Avg. products per market	279	242	355
Nr. firms	55	48	12
Avg. firms per market	3.25	3.06	3.65

Notes. The variables *Fare*, *Nr. Destinations*, and *Nr. Flights*, are in units of 100, and *Nr. Flights to Destination* is in units of 10, *Spacing* is in average hours between adjacent flights on the same day, HHI is on a scale of 0-1, *Connectivity* is in 100 flights arriving or departing from-to the connectivity window, *Fuel Consumption* is in 100 €per seat. Other variables are in units or dummy variables.

Sieg (2023)). Each under-filled short-haul flight represents a slot that could serve a higher-value long-haul route or be reallocated to a more efficient carrier. This pattern reverses on longer routes, where Air France achieves higher load factors than competitors (86.9% vs 84.3%), suggesting the dominant carrier strategically concentrates efficient operations on high-value routes while maintaining excess capacity on short routes as a competitive barrier.

Air France's slot advantage extends beyond quantity to quality: 26% of its flights depart in the morning versus just 18% for competitors, who concentrate operations at the end of the day. Morning slots are arguably more valuable, enabling full-day travel and providing operational buffers against cascading delays. Asymmetries are maintained by the historical accumulation of the most desirable slots, reinforcing barriers to competition.

Finally, slot dominance links to pricing power. Where Air France holds a high slot share, its fares are approximately 20% higher (€5.31 vs €4.42), while competitors' fares show no significant variation across these market types (€3.86 vs €3.96, $p = 0.58$). If high-slot-share markets were intrinsically more valuable — due to higher demand, wealthier passengers, or premium destinations — all carriers would charge higher fares there. The fact that only Air France benefits suggests the premium reflects market power derived from slot dominance rather than underlying demand characteristics. This is reinforced by the structural differences across markets: those with high Air France slot shares have fewer competitors (1.84 vs 2.25 firms) and higher concentration (HHI of 0.70 vs 0.58).

However, plausible alternative explanations exist. Air France's high frequency on short-haul routes may reflect business travelers' willingness to pay for schedule flexibility, the network value of feeding passengers into long-haul connections at its CDG hub, or cost advantages that make frequent service privately optimal. The raw statistics cannot distinguish strategic inefficiency from rational responses to demand and cost fundamentals, and, moreover, the flights and network operated by Air France differs from other firms on many observable dimensions (Table 3). The structural model addresses this by separately estimating demand for schedule characteristics, the network value of connectivity, and carrier-specific marginal costs as a function of network observables.

II Market equilibrium conditional on slot allocation

A departure (or arrival) slot gives the carrier the right to schedule one flight departing from (or arriving at) that airport at that specified time and day of the week. This flight generates profit for the operating carrier that depends, 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.

II.A Demand

II.A.1 Setup

The demand model is based on Berry, Carnall and Spiller (2006) and Berry and Jia (2010). Two latent consumer types vary in their preferences for product characteristics, including 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 ($j = 0$). The indirect utility for consumer i of type r for consuming flight j with non-price characteristics X operated by carrier a to market t is:

$$U_{ijt}(\mathcal{F}_{at}) = \alpha_r P_{jt} + N_{jt}^D(\mathcal{F}_{at})'\beta_r^N + X_{jt}'\beta_r + \mu_t + \mu_a + \epsilon_{ijt}(\lambda), \tag{1}$$

where $U_{i0t} = \lambda\zeta_{i0t}$ for the outside option of not flying, and μ_a and μ_t are carrier and destination fixed effects, the latter also captures alternative travel modes such as high-speed trains. The composite error term $\epsilon_{ijt}(\lambda)$ equals $\xi_{jt}+\nu_{it}(\lambda)+\lambda\zeta_{ijt}$; ξ_{jt} captures the flight-level scalar unobserved to the econometrician but (potentially) observed by firms and consumers; ν_{it} captures consumers' preference for flying over not flying to market t; and ζ_{ijt} captures consumer i's idiosyncratic utility for flight j to market t. Parameter $\lambda \in (0, 1]$ governs how closely substitutable flights are relative to the outside option. If $\lambda = 1$, the model reduces to a simple logit where substitution to the outside option is as likely as substitution to another flight within the same market.

The notation $U_{ijt}(\mathcal{F}_{at})$ indicates that demand for a flight depends on how it is positioned in the carrier's network, as measured by the demand-relevant network characteristics $N_{jt}^D(\mathcal{F}_{at})$, such as the carrier's airport presence. X_{jt} is a vector of flight-level characteristics that consumers value independently from the network, such as departure time. A full description of which network variables are included and why they are plausibly exogenous is provided below. Equilibrium prices are potentially endogenous and will be instrumented. Furthermore, ζ_{ijt} is assumed Gumbel-distributed, and ν_{it} is drawn from the distribution of Cardell (1997), which ensures $\epsilon_{ijt}(\lambda)$ is Gumbel-distributed and yields the familiar closed-form nested logit purchase probabilities.[9]

[9] The market share of product j in market t is:

$$S_{jt}(\mathcal{F}_t; \xi_t, \theta^D) = \gamma_1 D_{j1} \frac{\left(\sum_{k=1}^J D_{k1}\right)^{\lambda-1}}{\left(1 + \sum_{k=1}^J D_{k1}\right)^{\lambda}} + (1-\gamma_1) D_{j2} \frac{\left(\sum_{k=1}^J D_{k2}\right)^{\lambda-1}}{\left(1 + \sum_{k=1}^J D_{k2}\right)^{\lambda}}, \tag{2}$$

where

$$D_{jr} = \exp\left((\alpha_r P_{jt} + N_{jt}^D(\mathcal{F}_{at})'\beta_r^N + X_{jt}'\beta_r + \xi_{jt} + \mu_a + \mu_t)/\lambda\right).$$

II.A.2 Flight-level demand with network effects

An operating carrier must have arrival and departure rights for each slot-constrained airport on the itineraries of its consumers, for which values must be isolated. Slot values are defined formally in Section IV, but in short, the carrier's willingness to pay for a slot is derived from the additional profits generated (across the carrier's network) with the single flight that can be scheduled in that slot, accounting for how this flight's position in the carrier's flight network affects demand and costs across its network. As such, the goal is to incorporate all major factors driving an airline's profits that can reasonably be ascribed to the flight in question, including externalities on its other flights.

The resulting flight-level analysis contrasts with the more typical itinerary-based approach (see also subplots a) and b) in Figure 1). An *itinerary* is a set of flights established between origin and destination 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%) that departed from CDG and none that departed from ORY in S18 were identified as *connecting flights*. In contrast, *Aéroports de Paris* has documented that 21.7% of the passengers departing from its airports (CDG and ORY) were connecting from another flight that year.[10] This means that a non-negligible portion of demand for (direct) flights comes from passengers who are on multi-legged itineraries.

To address this, 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 airport and the destination airport. Specifically, flight characteristics include the number of other feasible flights by the same carrier or a codesharing partner that arrives in (depart from) the departure airport (arrival airport) within a time window of 1–6 hours before the flight takes off (after the flight arrives). The feasibility constraint imposes that a flight in these arrival or departure windows is only counted as a potential connection if the flight distance with the connection in Paris would be less than 1.5 times the direct flight distance between the origin and destination of the itinerary (as in Yuan and Barwick (2024)). This rules out that consumers consider flying through Paris to travel between two relatively nearby places.[11] The connectivity variables reflect the convenience of the flight

γ_1 is the population share of type 1, and $\theta^D = (\lambda, \gamma_1, \alpha_1, \alpha_2, \beta^N, \beta)$ is the 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. 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 ($\mathcal{F}_t$).

[10] Source: press release *Aéroports de Paris* January 15 2019. Overall, the global schedules data contains the universe of 32,342,819 flights for the 2018 summer season, of which 26,800,445 were scheduled passenger flights considered relevant for this study. A total of 86.6% of all scheduled passenger flights globally are direct, 10.5% have one stop, and the remainder have two or more stops. The 128 flights with connections at CDG were for only 6 routes, as documented Marra (2024).

[11] As there is a sizable offering of flights from low-cost carriers with point-to-point network strategies in Europe, it is not mandated that the connection airport is a hub, and simple connectivity measures without the feasibility constraint are also considered.

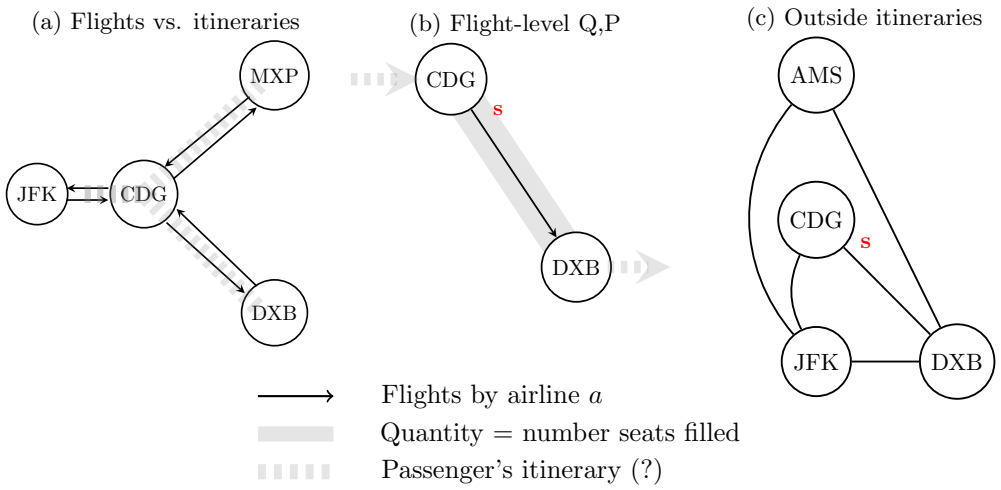


Figure 1: Flight-level analysis to recover the value of slot **s** to airline a

Notes. From the perspective of a consumer, JFK-CDG-MXP and JFK-CDG-DXB are two different itineraries both with a connecting stop at CDG (plot a), and modelled typically as two non-directional or four directional markets. From the perspective of an airline carrier, three flights are involved in each direction. Slots are required for all slot-constrained airports on potential itineraries. In this paper, slot values are therefore estimated off the demand for the (direct) flight associated with the slot (**s** in plot b) incorporating externalities on the rest of the flight network — including, through network (opportunity) costs, potential outside itineraries (e.g., JFK-AMS-DBX in plot c) and hub functions.

in relation to other flights on a (potential and unobserved) itinerary, which itinerary-level analyses account for with, e.g., the number of connections (Berry and Jia, 2010) or the ratio of itinerary distance over nonstop distance (Gayle, 2013). Estimating high connectivity premia will make a slot more valuable to airlines who can use it to offer an itinerary with a short layover at CDG or a destination airport.

Connecting passengers are also accounted for in the market (size) definition.[12] Markets are defined as cities that can be reached in a given month by nonstop flight from Paris, conform the directional city-pair market definition and with the estimation sample containing only one airport per destination city. CDG and ORY airport are substitutes in demand.

As a result, I recover slot-related consumer preferences. Most importantly, consumers likely have departure and arrival time preferences. Such time preferences will cause some slots to be more valuable than others, as airlines enjoy higher mark-ups in equilibrium for flights scheduled at more favorable times. Without this feature, it will be hard to capture that departure slots are more desirable i) during peak times or ii) when they are suitable

[12] The market size equals the geometric mean of the population in the destination city and the population in Paris weighted by the overall share of nonstop passengers in Paris (0.873), plus the average geometric population mean across all destination cities weighted by the share of connecting passengers (1-0.873). This reflects that a nonnegligible share of demand for flights departing from Paris is by passengers flying from or to cities beyond the origin or destination of the flight.

to arrive at favorable times at profitable long-haul destinations.[13] Time preference will therefore be instrumental to help predict instances like the record price of $75 million that Oman Air paid in 2016 to Kenya Airways for a single slot pair at Heathrow Airport. Oman uses the slot for an overnight flight from Muscat to Heathrow, and needs the early arrival time to cater to businessmen that make further connections to the US and Europe.[14] Likewise, day of the week preferences are included in demand and consumers may prefer one departure airport over another due to the implied time it takes to reach the airport and then the gate.

Further, the carrier's airline presence captured by the number of destinations it serves is an important demand driver (as documented in the literature dating back at least to Borenstein (1989), Berry (1990), and Brueckner, Dyer and Spiller (1992)). The positive airport presence effect reflects numerous advantages, such as having access to better lounges or more convenient gates. Carrier (group) fixed effects are included in the utility specification, in addition to network variables, to capture this enhancement. In addition, more flights on the same date to the same destination by the same carrier might be preferred if this reduces the expected wait time when a flight is canceled. The utility specification in (1) contains:

$$N^D = \left[ConnectivityOrigin, ConnectivityDest, \log(NrFlightsToDest), NrDestinations\right]$$
$$X = \left[DepartureAirport, DepTime, DepDay, ArrivalTime\right]$$

The specification serves to capture consumer preferences in relation to a carrier's network of flights and the role that a certain departure slot could have in it. Including the described network-based variables in the demand model makes it unclear whether consumers are better off when transferring a slot from the dominant carrier to a competitor in the same market, even when prices decrease. The estimated model will illuminate a fundamental trade-off between higher prices and greater convenience associated with concentration in this industry.

II.B Marginal costs

The marginal cost per seat of flight j to market t is given by

$$C_{jt}(\mathcal{F}_a) = N^S_{jt}(\mathcal{F}_a)'\kappa^N + W'_{jt}\kappa + \omega_{jt}, \tag{3}$$

where $N^S_{jt}(\mathcal{F}_a)$ contains network-related variables such as the average spacing between departures, W_{jt} captures flight-level cost shifters such as the aircraft's fuel consumption, $\theta^S = (\kappa^N, \kappa)$ is a vector of cost parameters, and ω_{jt} is an additive idiosyncratic cost term

[13] Given short turnover times, departing at peak times means arriving at peak times as well, as is relevant when thinking about matching arrival-departure slot pairs in the counterfactuals.

[14] As documented in this 2016 annual report by the Royal Aeronautical Society's Greener by Design body.

assumed to satisfy $(N_{jt}^S(\mathcal{F}_a), W_{jt}) \perp \omega_{jt}$. Marginal costs are taken to be independent of the number of flights on a passenger's itinerary, as it is unlikely that the sequence of flights a passenger takes influences the cost of operating them. Fuel expenditures constitute a large share of the variable costs of operating commercial aircraft. The average fuel consumption per seat is approximated by an aircraft-specific quadratic function in the distance to destination, using the aircraft-specific fuel efficiency coefficients estimated in Seymour et al. (2020), which mimics physics-based fuel burn performance models used in the industry.[15]

Carriers' flight networks naturally affect costs, mainly through the economies of density associated with hub-and-spoke relative to point-to-point networks. Quantifying these network-based cost differences is important for computing differences in (counterfactual) profits across airlines. Transporting passengers from or toward hubs is beneficial, as larger aircraft can be used and/or load factors increase. The cost specification includes the number of destinations and flights the carrier serves from the departure airport and from the arrival airport. Flight timing may also be relevant for costs: flights departing near each other can share check-in personnel, for example. Inspired by Borenstein and Netz (1999), I include the average spacing in hours of adjacent departures (by the same carrier on the same date to any destination). Flight spacing lends another time dimension to the model; estimating positive flight spacing costs and connectivity premia would endogenously generate a preference for slots that allow flights to be banked together (cf. the mechanism in Adler, Andreana and De Jong (2026)) rather than spaced out or randomly positioned.

I furthermore include a departure airport dummy to account for passenger and handling fee differences, whether the aircraft is wet-leased from another carrier, the share of business class seats, aircraft capacity, and flight distance bins. Arrival time variables further place the flight in the carrier's network, capturing the potential benefit of arriving in time for a new departure from the arrival airport at a favorable hour.

Taking economies of density and hub-and-spoke strategies seriously, the specified cost function is more accurately interpretable as the average variable cost per seat on flight j *net of any benefits* of transporting passengers to destination t — or equivalently, the opportunity cost of the seat in the short run when network decisions are already made. This interpretation is without loss if variable costs are linear in the number of passengers. Marginal costs are therefore not required to be strictly positive. Negative costs may arise, for instance, when fares are substantially discounted on feeder flights transporting passengers to

[15] Physics-based models simulate the geometric, kinematic, and kinetic properties of aircraft performance and require a wealth of data, such as cruising altitude, lift coefficient, climbing speed, and flight envelope data (Seymour et al. (2020)). See, for example, https://www.eurocontrol.int/model/bada for EUROCONTROL's aircraft performance model. Fuel consumption per seat is multiplied by the Kerosene-type Jet Fuel Spot Price one month before the departure date in euros to account (imperfectly) for input price variation over time. The prior literature typically uses flight distance; (Bauer (2025)) adopts my cost shifter as well.

a hub. For the same reason, the opportunity cost of flying to a slot-constrained destination is expected to be lower, since demand for onward departures from those airports will be higher. The marginal cost specification in (3) entails:

$$N^S = \big[NrDestinations, Spacing, NrFlights, NrDest(atDest), Spacing(atDest), \\ NrFlights(atDest), SlotConstrainedDest, ArrivalTime\big]$$
$$W = \big[FuelConsumption, WetLease, DepartureAirport, BusinessClassShare, \\ LongHaul, MediumHaul, Capacity\big]$$

This model generates spillovers across markets (as in Bontemps, Gualdani and Remmy (2023)) as well as across products. Consider an airline that receives a slot and adds a flight to a certain destination. If consumers value flight frequency, demand for *all* flights by this carrier to that destination increases, ceteris paribus. Moreover, the airline potentially enjoys cost savings from reduced spacing and a denser network, lowering the marginal cost of all flights to any destination on that date. These network effects are taken into account when computing slot values by updating $N^S(\mathcal{F})$ and $N^D(\mathcal{F})$ when the flight network $\mathcal{F}$ changes.

II.C Network (opportunity) costs

As the model is used to elicit ordinal preferences on slots, omitted fixed costs of operating flights only affect the results insofar as they vary across airlines above and beyond factors already included in the cost specification. Relatedly, as in most of the prior literature, I abstract from the decision to cancel a flight when realized demand is insufficient to cover fixed costs. The "use-it-or-lose-it" rule that applies at both slot-constrained departure airports and the majority of destinations in the sample makes cancellation particularly unattractive. However, the different aircraft fleets and global slot holdings of carriers constrain them in using new slots optimally. Moreover, carriers with extensive outside hub networks face a higher threshold for operating Paris flights: passengers can be routed through alternative hubs, so a Paris slot is only worthwhile if direct-route profitability is sufficiently high. Conversely, carriers dependent on Paris as their primary hub will tolerate lower-margin flights to maintain connectivity and the model estimates are used to capture this (Section III.C).

II.D Equilibrium

Multiproduct firms compete in average flight fares for consumers within a market, given externalities across their flight network as measured by the network variables in the cost and utility specifications. Demand is assumed to be independent across markets, and equilibrium

prices are obtained market-by-market. Assuming that observed prices maximize the firm's profits across markets, the equilibrium markups and quantities are recovered for each flight from the standard Bertrand-Nash first-order conditions. This delivers airline a's flight-level equilibrium profits $\mathbf{K}_a^*(\mathcal{F})$, which stacks the flight-level profits K_{jt}^* $\forall(j,t) \in \mathcal{F}_a$, and is used to determine slot values to carrier a below. The equilibrium can be interpreted through the lens of decentralized decision-making, similar to Aguirregabiria and Ho (2012). A local pricing manager prices all the carrier's flights out of Paris, disregarding how these prices affect demand for flights departing from other airports. Externalities on flights arriving in Paris or departing from the destination are internalized, but substitution to flights on outside itineraries not involving Paris is ignored.[16] Organizational frictions might justify abstracting from network externalities in pricing (Hortaçsu et al. (2024)), but model simulations defining slot values address that carriers' outside options affect the value they place on Paris slots.

II.E Identification and estimation

The identification of model parameters follows Berry and Jia (2010) and the results in Berry and Haile (2014) for nonparametric differentiated product models with market-level data and scalar unobserved heterogeneity. Identification requires at least one valid instrument for each endogenous variable, satisfying exclusion and completeness restrictions.

The price instruments (Z^D) must generate variation in average fares across flights in a market. The primary source is the fuel consumption per seat instrument based on fuel efficiency coefficients from Seymour et al. (2020) and flight distances, as described earlier. This instrument alone explains 79% of price variation — consistent with evidence that fuel costs constitute a substantial share of operating expenses (Brueckner, Kahn and Nickelsburg, 2024). The identifying assumption is that consumers don't value (or observe) the aircraft's fuel efficiency, which seems reasonable for the average traveler in 2018. Z^D also includes the number of products in the market, network-based instruments (shares of monopoly and duopoly markets at the destination), and BLP-style instruments based on flights' positions in characteristics space.[17] Instrument validity is assessed via a plain logit specification (Berry, 1994). Capacity and wet-lease instruments fail the Hansen J test under market-level clustering, consistent with airlines selecting aircraft on the basis of unobserved demand. Schedule-based instruments—such as the number of competing flights within a

[16] For a passenger travelling from New York to Dubai, flying through Paris is just one of several options: other European airports serve as major connecting hubs, and direct flights may be available — as illustrated in subplot c) of Figure 1. A more sophisticated pricing strategy would account for the fact that an airline controlling many alternative itineraries might exploit this by charging higher fares on the CDG–Dubai flight and being more selective about which Paris routes to operate. One approach would be to model this as a direct pricing externality across CDG–Dubai flights. Without demand data on outside flights, however, doing so is infeasible without substantial additional structure and restrictions on outside flight margins.

[17] Their identifying power may be reduced by large market asymptotics (Armstrong, 2016).

two-hour window, the number of seats on competing flights, and average flight spacing at the destination—pass this test, as expected given that slot constraints at congested airports impose external restrictions on timing that are plausibly orthogonal to demand shocks. The Wu–Hausman test confirms price endogeneity, and weak instrument tests confirm relevance.

The same instruments enter Z^S to identify the supply side, affecting quantities through equilibrium markups. Since aircraft characteristics are potentially endogenous — airlines may adjust them based on realized demand — they are instrumented with additional instruments exploiting airline fleet composition: the carrier's average capacity and small-plane share across other routes and rivals' capacity choices on the same route. These reflect long-run capital and network decisions plausibly orthogonal to route-specific demand shocks.

The model thus takes as given the aircraft fleet, route network, and flight scheduling, determined before slot auctions occur. At level-3 slot-constrained airports like CDG and ORY, the majority of seasonal slot allocations derive from historic rights.[18] While historic rights constrain new flight offerings, carriers might still reallocate slots in response to realized shocks. If slot usage changed frequently, estimation would need to address this selection (Draganska, Mazzeo and Seim, 2009; Fan and Yang, 2020; Ciliberto, Murry and Tamer, 2021; Li et al., 2022). However, the exogeneity of schedule-based instruments suggests that scheduling patterns do not respond to unobserved demand shocks c.p., supporting the assumption that slot rigidities generate exogenous variation in competitive pressure.

In the resulting *incremental* slot allocation model, the observed flight network is assumed independent of current-season shocks; the timing is visualized in Figure 2. The counterfactual results will therefore reflect primary auctions of small capacity shares — such as slots returned to the pool or relaxed constraints on secondary markets — rather than mechanisms allocating all capacity simultaneously. At the margin, changes to the slot allocation are simulated with current-period shocks observed even when the observed networks are independent of these shocks.

The demand parameters are obtained via a standard generalized method of moments (GMM) estimator, instrumenting for endogenous prices. The supply parameters are recovered with an IV regression of the resulting costs (prices minus equilibrium markups) on observables, instrumenting for endogenous quantities, aircraft capacity, and wet-lease status. Additional details on identification and estimation are provided in Appendix C.

[18] Pouget et al. (2023) document that fewer than 14% of slot requests at CDG in summer 2018 were for flights not offered in the previous equivalent season. This aligns with Milioti and Odoni (2024), who report that historic series account for at least 90% of allocations at Europe's busiest Level 3 airports.

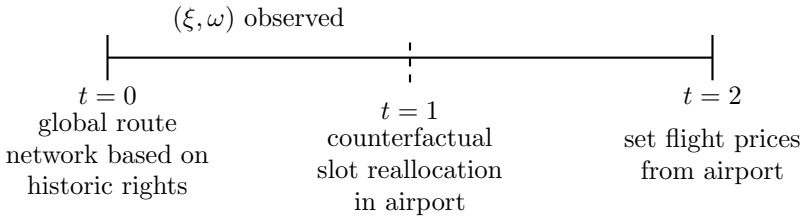


Figure 2: Incremental slot allocation model.

Notes. The global route network formation takes place at $t = 0$, before a small share of slots at one airport gets (re)-allocated in period $t = 1$. The route network formation depends on route-level fixed cost shocks that are independent of the idiosyncratic flight-level demand (ξ) and marginal cost (ω) shocks, which realize after the global flight networks set.

III Estimation results

III.A Demand parameters

The estimation results (Table 4) indicate that a small share of the population ($\hat{\gamma}_1 = 0.052$) has preferences diverging substantially from the majority. Based on the estimated price coefficients, Type 1 consumers exhibit lower price sensitivity ($\hat{\alpha}_1 = -0.445$) and can be labeled as business travelers. Type 2 consumers, comprising all other passengers including leisure travelers, are substantially more price-sensitive ($\hat{\alpha}_2 = -1.462$).[19]

The nested logit coefficient ($\hat{\lambda} = 0.787$) is significantly different from one, confirming that the nesting structure captures meaningful substitution patterns, and is similar in magnitude to prior studies. The price elasticity of the leisure type (-5.58) is comparable to the value found in Berry and Jia (2010) (-5.0 in the base specification), while the smaller share of business types is less price elastic (-1.69 vs. -0.44). The resulting aggregate price elasticity is twice as large (-3.48 vs. -1.55). Note that the price elasticities in most airline studies (including Berry and Jia (2010)) are based on aggregated data, which likely introduces an upward bias as documented in D'Haultfœuille et al. (2022) in the context of rail transportation. My estimated aggregate price elasticity of -3.48 is close to the -3.31 average found in Williams (2022), also estimated on flight-level data. The nesting parameter is stable across all specifications with destination fixed effects.

Consumers consistently rank Air France above Lufthansa and easyJet, but prefer other flag carriers (e.g., Singapore Airlines, Turkish Airlines) and non-state-owned carriers (e.g., Delta Air Lines, United Airlines) over Air France. Table B1 lists the carriers by category. To rule out that these rankings primarily reflect differences in destinations served rather than true brand preferences, I estimate the model excluding monopoly routes. The results in

[19] Because the more price-sensitive leisure types are roughly four times more likely to choose the outside option (e.g., taking a train, forgoing the trip altogether), the share of each type actually on board is more evenly balanced — with business travelers filling around 20% of seats on the plane on average across markets.

Table 4: Demand estimation results

		Alternative estimation samples		
		No Overseas	No Monopoly	No CDG
	(1)	(2)	(3)	(4)
Model parameters				
Nested logit coefficient (lambda)	0.787 (0.008)	0.785 (0.008)	0.782 (0.010)	0.742 (0.010)
Population share type 1	0.052 (0.003)	0.053 (0.003)	0.052 (0.004)	0.125 (0.013)
Fare (type 1)	-0.445 (0.015)	-0.420 (0.015)	-0.474 (0.021)	-0.517 (0.020)
Fare (type 2)	-1.462 (0.022)	-1.487 (0.025)	-1.454 (0.023)	-1.419 (0.052)
Network & time preferences				
Departure Airport = CDG (type 1)	-0.596 (0.107)	-0.706 (0.105)	-0.593 (0.125)	
Departure Airport = CDG (type 2)	0.437 (0.036)	0.445 (0.034)	0.441 (0.038)	
Arrival time cosine (type 1)	0.646 (0.032)	0.630 (0.030)	0.526 (0.050)	1.180 (0.103)
Arrival time cosine (type 2)	0.019 (0.017)	0.014 (0.017)	0.033 (0.020)	0.528 (0.055)
Nr. Destinations	-0.002 (0.024)	0.042 (0.023)	-0.003 (0.026)	-2.389 (0.229)
log(Nr. Flights to Destination)	0.071 (0.012)	0.040 (0.012)	0.069 (0.012)	0.079 (0.016)
Connectivity Arrival Airport	0.115 (0.010)	0.112 (0.010)	0.118 (0.011)	0.142 (0.016)
Connectivity Departure Airport	0.344 (0.012)	0.339 (0.012)	0.359 (0.012)	0.108 (0.043)
Day of week fixed effects				
Day = Tuesday	0.154 (0.013)	0.155 (0.012)	0.195 (0.014)	-0.075 (0.013)
Day = Wednesday	-0.247 (0.011)	-0.237 (0.011)	-0.243 (0.012)	-0.155 (0.014)
Day = Thursday	0.198 (0.013)	0.203 (0.012)	0.240 (0.014)	0.052 (0.015)
Day = Friday	-0.106 (0.012)	-0.095 (0.012)	-0.103 (0.013)	0.059 (0.016)
Day = Saturday	-0.038 (0.013)	-0.032 (0.013)	-0.023 (0.014)	0.101 (0.018)
Day = Sunday	0.026 (0.013)	0.033 (0.012)	0.032 (0.013)	0.066 (0.016)
Carrier fixed effects				
Air France	-0.108 (0.038)	-0.264 (0.036)	-0.132 (0.039)	0.667 (0.097)
easyJet	-0.398 (0.034)	-0.546 (0.032)	-0.406 (0.034)	0.266 (0.053)
Other Flag Carrier	0.148 (0.041)	0.062 (0.042)	0.154 (0.042)	0.080 (0.054)
Lufthansa	-0.425 (0.057)	-0.482 (0.057)	-0.446 (0.059)	
Vueling Airlines	-0.073 (0.037)	-0.209 (0.037)	-0.077 (0.038)	0.109 (0.054)
Observations	119,810	117,217	111,534	38,931
MMSC-BIC (scaled by N)	-0.001	-0.001	-0.001	-0.003
Share inelastic demand	0.001	0.001	0.001	0.029
Avg. own-price elasticity	-3.476	-3.385	-3.568	-2.436
Avg. Lerner index	0.424	0.429	0.412	0.646

Notes. Estimated demand-side coefficients with standard errors clustered at the market level in parentheses, for the full estimation sample (column 1) and robustness checks excluding French overseas territories (column 2), monopoly destinations (column 3), and all flights departing from CDG (column 4). All specifications pass Hansen J-tests ($p > 0.99$) under destination-level clustering. Lufthansa only operates out of CDG. Destination fixed effects are omitted; Appendix Table B3 lists the 10 most and least attractive destinations as estimated in the full sample. Units are as described in the footnote of Table 3

column 3 confirm that the carrier ranking is unchanged and other estimates are largely stable, supporting the interpretation that these coefficients reflect genuine consumer preferences rather than supply-side market structure.

The estimated destination fixed effects capture the utility of flying to each destination relative to alternative modes such as rail. As expected, destinations with the lowest flight utility are those accessible by high-speed train, including London, Milan, and Marseille, while those with the highest utility — Atlanta, Saint-Denis, and Dzaoudzi — are separated

Table 5: Price elasticities, willingness to pay, and diversion patterns

Own price elasticities	
Type 1	-1.692
Type 2	-5.580
Both types (population weighted avg.)	-3.476
Share of passengers	
Type 1	0.052
WTP departing from CDG vs ORY	
Type 1	-133.982
Type 2	29.888
WTP +100 potential connections	
Departure Airport - Type 1	77.294
Departure Airport - Type 2	23.510
Arrival Airport - Type 1	25.788
Arrival Airport - Type 2	7.844
WTP arriving at 7am vs 1pm	
Type 1	145.378
Type 2	1.302

	Diversion	
	Same carrier	Rival
Comparing haul type (AF CDG, 06–10h)		
Long-haul	74%	26%
Short-haul	60%	40%
Comparing airport/destination set (AF long-haul)		
AF ORY	34%	66%
AF CDG	74%	26%
Price spillover AF CDG long-haul (peak):		
Same time window	+0.01% of fare	
Adjacent time windows	+0.47% of fare	
Comparing time windows (AF CDG long-haul)		
Morning (06–10h)	74%	26%
Midday (10–14h)	63%	37%
Afternoon (14–18h)	82%	18%

Notes. Left panel: willingness to pay reported in euros per ticket; type 1 can be labeled a "business" type and type 2 as traveling for "leisure". Right panel: diversion shares conditional on purchasing. AF = Air France. Price effects are average fare changes on same-carrier flights in the same vs other 2-hour departure windows. AF CDG long-haul averages over all time blocks.

from Paris by an ocean. The sample includes five French overseas territories (with destination cities Saint-Denis, Dzaoudzi, Cayenne, Fort-de-France, and Pointe-à-Pitre), which are subject to specific subsidy and tax schemes. Estimating the model excluding them (column 2) leaves most estimates unchanged, although the Air France coefficient becomes more negative, suggesting travelers to these destinations have an above-average preference for the former state-owned carrier.

The estimates reveal sizeable departure and arrival time preferences. Most strikingly, business travelers are willing to pay €145 more to arrive at their preferred morning time — in time for their business day or a connecting flight (Table 5). At about 150% of median profits per seat, this rejects the hypothesis that airport time slots are perfect substitutes and explains the large sums paid to secure early morning slots.[20]

Diversion patterns suggest that destination reachability is an important factor differentiating slots at different times (Table 5, right panel). Removing an Air France long-haul slot in the morning at CDG retains 74% of within-market demand within the carrier versus 60% for short-haul at the same airport and time, because long-haul routes face less competition at the same time. The same long-haul morning slot at ORY retains only 34%, reflecting a more competitive destination mix. Price spillovers reinforce this: the fare effect concentrates

[20] Appendix D compares alternative specifications, also highlighting an undesirable consequence of restricting passenger types to have identical time preferences. Under such restrictions, business and leisure travelers are predicted to substitute in roughly equal proportion to other morning flights (driven only by price differences). Models with type-specific time preferences obtain more realistic differentiated diversion: business types divert 2.9 times as much as leisure types to other morning flights. This makes both the size and timing of slot holdings relevant for market power.

in adjacent time windows (+0.47%) rather than the same window (+0.01%), as displaced passengers rebook into the next long-haul bank. Within Air France's CDG long-haul flights, within-carrier diversion falls to 63% at midday when US carriers compete on transatlantic routes, versus 82% in the afternoon bank when Air France faces almost no competition.

The estimated connectivity coefficients indicate that travelers value shorter layovers. Business and leisure travelers are willing to pay €77 and €24, respectively, for 100 additional (potentially connecting) flights arriving at the departure airport within 1–6 hours before their flight — about one standard deviation more connectivity. Willingness to pay for departure airport connectivity is roughly three times that for arrival airport connectivity, perhaps because arrival time preferences already partially capture this value. Reallocating slots from larger to smaller networks thus reduces consumer convenience from short layovers, and implies a lower connectivity premium for the smaller-network carrier.

Consumers prefer to fly with a large carrier from the departure airport, though this preference is not explained by the number of destinations served. Instead, consumers value carriers offering more same-day flights to their desired destination, which enters utility in logs to reflect decreasing marginal gains from higher frequency. Business travelers also strongly favor ORY and are willing to pay €134 to avoid CDG, which is less convenient from central Paris and notorious for inter-terminal transit times.[21]

Estimation excluding flights from CDG (column 4) reveals furthermore that business travelers comprise a larger share at ORY (12.5% vs 5.2%), consistent with time-sensitive travelers valuing ORY's proximity to central Paris. Connectivity at the departure airport matters less at point-to-point ORY than at hub CDG, where passengers value access to Air France's connecting banks. The larger arrival time coefficient for both types and similar price sensitivity further suggest that time is more relevant for passengers departing from ORY. With a different mix of legacy and low-cost carriers at ORY (Table B1), Air France becomes the preferred carrier out of this airport. Despite these differences, counterfactuals are based on the main specification estimated jointly on both airports: they serve overlapping catchment areas with substitution across airports, and slot reallocation at one airport may shift passengers and competitors to the other. Appendix D provides additional details.

III.B Supply parameters

The cost estimates show that the flight network generates cost synergies (Table 6). I estimate diseconomies of scope — a larger network is more costly for a given number of flights —

[21] If CDG and ORY served non-overlapping destinations, a CDG utility premium could reflect destination preferences as much as airport preferences. This is not the case: of 119 destinations, 48 are reachable from both airports, 8 exclusively from ORY, and 63 exclusively from CDG. The counterfactual simulations furthermore select only from carriers' existing flight networks, so that even if the CDG effect captures some mix of airport and destination benefits, the same mix drives welfare in the counterfactuals.

Table 6: Marginal cost estimates

	(1)	Shadow Cost (75%) (2)	Shadow Cost (90%) (3)
Constant	0.321 (0.019)	0.389 (0.023)	0.416 (0.025)
Nr. Destinations	0.726 (0.016)	0.698 (0.017)	0.681 (0.017)
Spacing	0.067 (0.002)	0.076 (0.003)	0.077 (0.003)
Nr. Flights	-0.238 (0.012)	-0.199 (0.014)	-0.186 (0.014)
Nr. Destinations (at Destination)	0.121 (0.007)	0.113 (0.007)	0.113 (0.007)
Spacing (at Destination)	0.020 (0.001)	0.021 (0.001)	0.021 (0.001)
Nr. Flights (at Destination)	0.131 (0.004)	0.138 (0.004)	0.138 (0.004)
Fuel Consumption	0.326 (0.004)	0.313 (0.005)	0.309 (0.005)
Wet Lease	-0.236 (0.032)	-0.276 (0.033)	-0.297 (0.034)
Departure Airport = CDG	0.175 (0.011)	0.167 (0.011)	0.175 (0.011)
Slot Constrained Destination	-0.237 (0.009)	-0.240 (0.009)	-0.234 (0.009)
Business Class percentage	0.015 (0.001)	0.016 (0.001)	0.016 (0.001)
Flight = Medium Haul	2.179 (0.018)	2.206 (0.019)	2.208 (0.019)
Flight = Long Haul	2.218 (0.027)	2.268 (0.029)	2.292 (0.030)
Arrival = Mid Day	0.288 (0.010)	0.274 (0.010)	0.277 (0.010)
Arrival = End Day	0.180 (0.009)	0.168 (0.009)	0.169 (0.009)
Capacity	-0.002 (0.000)	-0.003 (0.000)	-0.003 (0.000)
Shadow Cost		0.023 (0.004)	0.210 (0.034)
Degrees of Freedom	119,793	119,792	119,792
Adj. R^2	0.8445	0.8450	0.8441
F-value vs. (1)		176.36	229.40
$p < 0.00001$		✓	✓

Notes. Estimated costs in 100 euros per ticket. Standard errors are clustered at the market level and reported in parentheses. Units are as described in the footnote of Table 3.

alongside economies of scale in the number of flights, and scheduling efficiency gains from bunching flights more closely together. Standardized effects of the cost shifters are given in panel (b) of Table 7. For example, a one-standard-deviation increase in flight spacing from the departure airport (by the same carrier and date, to any destination) increases costs by 8%, holding other variables at their sample means.

As expected, flight distance and fuel consumption are the most important cost drivers. Flights from CDG cost approximately €18 more per ticket partly because the supplemental per-passenger fee at CDG is more than twice that at ORY, and connecting passengers at CDG are more expensive.[22] Flights with a higher business class share are also more costly to operate. Arriving in the middle or end of the day is estimated to be costly, which may reflect the network benefit of having the aircraft available for an onward flight at a popular hour. For 1.5% of flights I estimate negative costs, interpretable as feeder flights where the benefit of routing passengers through the destination outweighs operating costs, or slot hoarding. The estimates also indicate that whether the destination airport is slot-constrained is a meaningful cost-side network variable: the net cost per seat is approximately €24 lower for

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

flights to slot-constrained airports, interpretable as the benefit of the aircraft reaching a popular destination. These results underscore that the estimated marginal cost should be interpreted as the average variable cost per seat *net of any direct network benefits*.

A potentially important dimension thus far omitted from the cost model is the opportunity cost of selling a ticket. The shadow costs associated with a fixed plane size may be economically significant, and I experiment with the inclusion of a soft capacity constraint. A soft capacity constraint reflects the common airline practice of *overbooking*; airlines are known to sell between 10% and 15% more tickets than seats (see Suzuki (2006) and the references therein). To capture this, I consider a smooth (sigmoid) approximation of the "hockey stick" cost of exceeding a given share of capacity outlined in Ryan (2012), where shadow costs begin when quantity exceeds $\varphi = 0.75$ or $\varphi = 0.90$ of the plane's capacity.[23] Shadow costs are a function of quantity and are treated as endogenous in estimation. Both specifications provide a better statistical fit than the unconstrained baseline ($p < 0.00001$, columns (2) and (3) in Table 6), indicating that capacity constraints play a meaningful role in explaining observed pricing, with the $\varphi = 0.9$ specification fitting best. While the estimates imply economically meaningful shadow costs (€37–131 per ticket when selling 15% more tickets than seats), they are not prohibitive, and the estimates are likely muted by observing occupied seats rather than sold tickets as quantity. Model simulations are therefore based on the cost specification without a shadow cost variable, but with a market-specific overbooking cost calibrated so that counterfactual equilibrium quantities never exceed 15% of capacity (in line with industry practice; Suzuki (2006)), to account for capacity constraints and avoid excessive overbooking on higher-margin flights.

Flights generate on average €20k in profits before fixed costs (Table 7). Profits are higher for long-haul flights, with substantial heterogeneity across distance categories. The median profit per seat is €55, the average is €98, and at the 90th percentile each seat generates €235 before fixed costs. Air France has similar profit statistics overall, with the exception of long-haul flights (higher: €66k vs. €57k per flight on average) and medium-haul flights (lower: €35k vs. €40k). I estimate an average Lerner index of 0.42, which is lower than the estimate for direct flights in Berry and Jia (2010), consistent with the higher price elasticity from the flight-level analysis and the richer cost specification, and reflects relatively tight margins in this industry.[24] The Lerner index for Air France flights is 0.43, close to their

[23] The soft capacity constraint adds:

$$\texttt{Shadow Cost}\ (\varphi, Q) = \exp(-(Q - \varphi\, Capacity))^{-1}\, \frac{(Q - \varphi\, Capacity)^2}{Capacity}, \tag{4}$$

where costs rise sharply once the load factor exceeds φ. The sigmoid form yields analytic standard errors and keeps marginal costs bounded, maintaining a well-defined GMM objective. Lazarev (2013), Williams (2022), and Aryal, Murry and Williams (2024) develop alternative estimators using hard capacity constraints in dynamic pricing models.

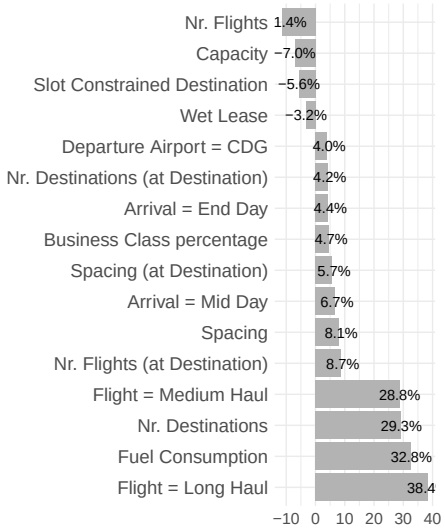
[24] Other estimates in the literature range from 0.32 to 0.39 — lower for direct flights without random

Table 7: Cost and profit estimates

(a) Estimated profits

	All	Air France
Lerner Index		
Average	0.42	0.43
Median	0.36	0.34
Avg. profits per flight (10k)		
All flights	1.98	2.04
Flights from CDG	1.17	1.18
Flights from ORY	1.18	1.23
Short Haul flights	1.10	1.12
Medium Haul flights	3.98	3.47
Long Haul flights	5.66	6.61
90th pct. profits per flight (10k)		
Short Haul flights	1.51	1.56
Medium Haul flights	6.97	6.78
Long Haul flights	9.40	10.96
Profits per seat		
Average	98.32	101.15
Median	55.10	56.76
90th percentile	235.25	235.25

(b) Cost estimates: marginal effects



Notes. The marginal effects in (b) give the percentage change in marginal costs when increasing the variable by one standard deviation, holding the other variables at their means. All estimated coefficients, given in Appendix Table 6, are statistically significant at at least the 5 percent level.

Table 8: Marginal Flight Profitability by Carrier

	Mean	Median	5^{th} pct.	95^{th} pct.	Air France
Profit of marginal flight	8146.69	5051.83	999.60	26443.94	**969.45**
Carriers with higher marginal profit than Air France: 52 of 55					
Bottom 5 carriers: Air Algérie, Air Caraïbes, **Air France**, Air Corsica, Aigle Azur					

Notes. Summary statistics of the flights with the lowest estimated variable profits of each carrier. The profit statistics give the mean, median, 5th and 95th percentile across the 55 marginal flights.

self-reported average gross profit margin of 48% across all Air France flights in 2018.[25]

Comparing the least profitable flights across carriers — their "marginal" flights — provides additional insight into the potential inefficiency of the current slot allocation, complementing the model-free assessment in Section I.C. The idea is that slot reallocation could improve welfare if historic grandfather rights allow the dominant carrier to retain slots despite low-value uses. The estimates support this: Air France operates flights with substantially lower marginal profits than other carriers (Table 8), with profits on its marginal flight eight times smaller than the average across the 55 carriers in the sample. Only Air

coefficients (Bontemps, Gualdani and Remmy (2023)), higher for integrated codeshare products (Gayle (2013)) — all based on U.S. markets using disaggregated DB1B data.

[25] Air France's profit numbers include profits from full freighters and from cargo in belly capacity, which I abstract from, see: Macrotrends.net and Air France's 2018 financial statement for further details.

Table 9: Estimated network opportunity costs

Panel A: Summary of C_{ah}			**Panel B: Determinants of C_{ah}**				
				4 carriers		55 carriers	
	Mean	Std. Dev.		(1)	(2)	(3)	(4)
Air France	18111	12541	Outside hubs ($HatO$)	914.07	-634.39	216.66	112.36
				(297.77)	(319.72)	(263.93)	(274.96)
Lufthansa	6733	1368	Outside flights ($JatO/1000$)	-772.65	168.86	-466.36	-701.70
				(164.86)	(252.12)	(275.37)	(267.79)
easyJet	6971	2653	Outside capacity (meanCap)	-42.92	-188.87	36.34	57.19
				(25.60)	(63.57)	(18.23)	(25.24)
Vueling	6166	1808	Paris flights ($JatF$)	128.23	-19.80	12.56	31.25
				(17.31)	(19.79)	(15.02)	(16.77)
All 4	9495	8109	Observations	72	60	303	128
All 55	14242	16341	Level	Carrier	Car.-hour	Carrier	Car.-hour
			R^2	0.789	0.487	0.142	0.250

Notes. C_{ah} is obtained by carrier-hour according to (5) given model estimates of demand and costs.

Algérie and Air Caraïbes exhibit lower marginal profits.

Even Aigle Azur — which went bankrupt shortly after the sample period — had 16% higher marginal flight profits. In this light, it is useful to assess whether the reallocation of the ORY slots returned by Aigle Azur was beneficial. Based on slot coordinator documentation, carriers receiving half of the slots under entrant status used them to add flights to Glasgow, Munich, Budapest, Sofia, Warsaw, and Frankfurt (see Appendix Table A3) while Aigle Azur operated flights to Agadir, Algiers, Bamako, Beirut, Berlin, Faro, Oran, and Porto in my estimation sample. I estimate that the consumer surplus was about 4 percent higher on average in the set of added markets. However, without knowing how carriers with non-entrant status used their slots and lacking details on the timing and characteristics of the newly added flights, this remains inconclusive.

An important caveat also applies to examining flights in isolation: network benefits accrue across flights, so a mass of low-margin flights could compound to substantial network value for both carrier and consumer (which the counterfactuals capture) and they might help sustain network value beyond Paris-destination pairs, which is addressed next.

III.C Network Opportunity Costs

Network opportunity costs are large when a carrier depends heavily on a hub and has few outside alternatives: losing a slot degrades a wide set of itineraries beyond the affected route, so the carrier is willing to operate individually marginal flights to preserve that presence. Conversely, a carrier with dense outside connections can reroute displaced passengers elsewhere, reducing the network cost of slot loss. If so, the variation in outside-network dependence (Table 2) should be visible in the cross-flight distribution of profitability: a carrier that retains slots for network reasons will operate some flights well below its own average,

whereas a carrier unconstrained by hub dependence will not. I capture this with C_{ah}, the gap between carrier a's average profit and its least profitable flight in hour h:

$$C_{ah} = \frac{1}{|\mathcal{F}_{ah}|} \sum_{f \in \mathcal{F}_{ah}} \hat{K}_f^* - \min_{f \in \mathcal{F}_{ah}} \hat{K}_f^* \tag{5}$$

where $\mathcal{F}_{ah}$ denotes all flights of carrier a departing in hour h and $\hat{K}_f^*$ the estimated equilibrium flight profit. A larger C_{ah} therefore reflects a wider spread between average and marginal flights, consistent with the carrier bearing higher costs of slot loss than direct-route profitability alone would justify.[26] Table 9 presents the estimates. Panel A reveals substantial cross-carrier heterogeneity consistent with this interpretation, focussing on the four carriers in the counterfactuals. Air France exhibits the largest values (€18k per slot-hour), reflecting its dependence on Paris as its primary hub. Lufthansa, easyJet, and Vueling show substantially smaller values (€6–7k), consistent with their greater ability to reroute passengers through outside hubs. Panel B regresses C_{ah} on outside network characteristics excluded by construction from the flight-level model. Larger outside flight networks are associated with lower C_{ah}, and greater Paris dependence with higher C_{ah}, in line with the proposed mechanism. The positive coefficient on outside hubs in column (1) is unexpected under hub-and-spoke logic but is consistent with the focal-city model of low-cost carriers, for which Paris substitutes for other bases rather than serving as a unique connection point. Column (2) introduces within-carrier variation by averaging network variables to the carrier-hour level; the correlation with outside hubs turns negative, consistent with hours serving better-connected destinations offering superior rerouting options. Columns (3)–(4) extend the sample to all 55 carriers, reducing explanatory power (R^2 falls from 0.789 to 0.142), though the negative association with outside flight networks remains statistically significant.

IV Slot values and welfare effects of market-based allocation

This section uses the model estimates to obtain slot values by simulating a carrier's total profits with and without a slot, depending on which flight it chooses to operate in the slot and what its competitors choose (the simulation method is summarized in Appendix F).

[26] The gap is more informative than the level because average profitability conflates network motives with route competition, costs, and demand. The dispersion within a carrier-hour cell instead reflects the coexistence of profitable and marginal flights, consistent with a hub-dependent carrier retaining slots it would otherwise drop. Route heterogeneity could generate a similar pattern, which is why the correlation with outside network characteristics provides validation.

IV.A Slots, carriers, and choice of flight

A slot is defined as a combination of departure airport, hour (between 6am and midnight), and day of the week. There are 245 unique combinations, and only one is "up for auction" sequentially, reflecting a directional approach that reallocates a small share of slots at an airport, or a secondary market where bilateral trades involve small network changes. The simulations focus on the four largest carriers operating from Paris: Air France, Lufthansa, easyJet, and Vueling Airlines, which together operate 60% of flights and sell 68% of tickets in the data. All other carriers hold a small market share individually; excluding them from the slot auction (but not the market) reduces the computational burden, and counterfactual expansion by smaller carriers would be more speculative. When selecting which flight to offer in slot s, airline a is restricted to flights from its existing Paris network ($\mathcal{F}_a$), updating only the timing of the flight and network variables when operating it in slot s. Restricting the set of carriers and their use of slots limits the risk of an exposure problem and reduces demands on the model to predict next-best destinations, which are not first-order for incremental slot changes. This approach is supported by the observed network changes in the schedules data: 92% of flights newly offered in summer 2019, by any carrier to any destination reachable from CDG or ORY, served a destination the carrier already flew to in summer 2018.

The value of a slot depends on how it will be used, but the rich flight-level model with spillovers in cost and demand makes a full optimal slot usage model computationally infeasible. The following assumptions reduce the computational burden.

The flight that airline a would operate when given slot s, $f_a(s)$, is taken to be the highest-profit potential flight by the airline, where potential flights are all observed flights by the carrier from the airport where the slot is located, subject to operational constraints. The most important restriction is that the destination must be served (by any carrier in the data) within a 3-hour window around the departure hour of slot s on any day, ensuring slots are matched with reasonable flights given time zones and flight times (e.g., *slot-destination feasibility*). I also impose that the same flight cannot be added to the schedule more than once. How each airline would use the slot is treated as public information, as can be known through the interaction of network managers at slot conferences.

The difference in flight offerings from Paris when airline a or airline a' owns the slot is assumed to consist only of the two flights they would operate in it, $f_a(s)$ and $f_{a'}(s)$, respectively. Moreover, since slots cannot be added to a congested system, a reallocation also implies that a flight is removed from the schedule ($f^X(s)$).[27] Throughout, it is maintained

[27] To determine $f^X(s)$, I first randomly draw a flight from the on average 488 flights per slot across 55 carriers, and identify the operator of that flight the "slot owner." $f^X(s)$ is then chosen as the flight generating the least baseline profit among all flights operated by the slot owner scheduled in s. As a robustness check, I also set $f^X(s)$ equal to the flight initially drawn.

that when one flight is added or removed from a carrier's network, the remaining network does not change in the short run. The binding slot constraints at many airports globally, and the complexity of network optimization under those constraints, support this assumption.

IV.B Allocative externalities and private values

Having established how each airline would use the slot, I now turn to what this implies for their slot values, defined as the difference in expected profits with and without the slot. Denoting the counterfactual flight network when carrier a obtains the slot by $\mathcal{F}^{s(a)}$, slot values are generally defined as:

$$\nu_a(s) = \sum \mathbf{K}_a^*\big(\mathcal{F}^{s(a)}\big) - \sum_{a' \neq a \in \mathcal{A}} \rho_{a'}(s) \Bigg(\sum \mathbf{K}_a^*\big(\mathcal{F}^{s(a')}\big) \Bigg) + C_{ah}(s), \tag{6}$$

$$\rho_{a'}(s) \in [0,1], \qquad \sum_{a' \neq a \in \mathcal{A}} \rho_{a'}(s) = 1, \qquad \rho_{a'}(s) \perp v_{a',s} \ \forall (a', a, s)$$

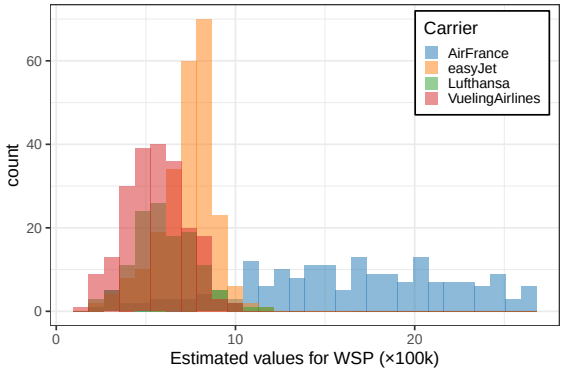
The first term captures profits when carrier a wins; the second is the expected profit when a competitor wins; the third adjusts for wider network value for carrier a at hour h.[28] The second term reflects that this setting features *allocative externalities*: equilibrium profits (denoted by the vector $\mathbf{K}_a^*(\mathcal{F})$, containing profits for all of carrier a's flights from Paris) depend on the full flight network. Carriers therefore care about which competitor obtains the slot if they do not, making this an *auction with aftermarket*. The alternative scenarios conditional on a losing are weighted by beliefs $\rho_{a'}(s)$ about which competitor would obtain the slot. Since the problem is intractable with fully general weights, I restrict $\rho_{a'}(s)$ to be independent of competitors' values. This restriction renders the setting one with *private values*: the realizations of competitors' cost and demand shocks do not affect $\nu_a(s)$, nor do competitors' valuations $\nu_{a'}(s)$. Private values also allow $\nu_a(s)$ to be estimated for each carrier separately and guarantee the existence of an efficient mechanism.[29]

The restrictions on $\rho_{a'}(s)$ rule out that firms incorporate competitors' win probabilities based on those competitors' $\nu_{a'}(s)$, but allow for a range of alternative heuristics. In the main specification, slot values are based on the firm's profits when winning versus when losing the slot to the competitor that would reduce their profits most:

$$\nu_a(s) = \sum \mathbf{K}_a^*\big(\mathcal{F}^{s(a)}\big) - \min_{a' \neq a \in \mathcal{A}} \Bigg(\sum \mathbf{K}_a^*\big(\mathcal{F}^{s(a')}\big) \Bigg) + C_{ah} \tag{7}$$

[28] The systematic correlation with observable outside network structure (Section III.C) support C_{ah} as a valid proxy for the hub-connectivity motives omitted from direct-route profitability. But as it is derived in parallel of the structural model, results are also presented excluding this measure.

[29] Jehiel and Moldovanu (2001) prove that when values are private, equilibrium existence in Bayes-Nash incentive-compatible mechanisms for multi-object auctions with interdependent values and allocative externalities is guaranteed.



Statistic	Value
Number slot values	861
Mean slot value (thousands)	855
Median slot value (thousands)	716
Mean network value share ($C_{ah}/\nu_a(s)$)	27.4%
Variance Decomposition:	
Across airports	4.7
Across days (same airport)	2.1
Across hours (same day-airport)	19.7
Across carriers (same hour-day-airport)	73.5

Figure 3: Estimated slot values

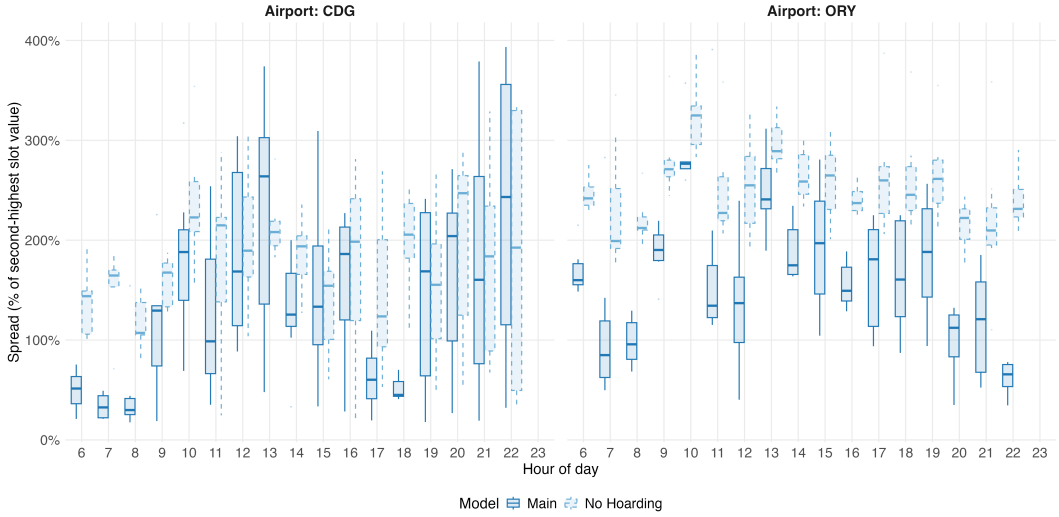


Figure 4: Spread between highest and second-highest slot values

Strategic slot hoarding is well-documented even among dominant carriers (e.g., De Wit and Burghouwt, 2008; Fukui, 2012; Fu et al., 2015), most vividly illustrated by Lufthansa CEO Carsten Spohr: *"We have to operate 18,000 additional, unnecessary flights during the winter, purely to secure our slots."* (Kösters, Meier and Sieg (2023)).[30] This behavior implies airlines place substantial value on denying slots to competitors or keeping them for future use, consistent with dynamic entry-deterrence motives (Aguirregabiria and Ho, 2012), and supports hoarding in (7); similar results obtain when firms ignore the opportunity cost of winning ($\rho = 0$).

[30]Ghost flights—near-empty aircraft operated solely to retain slot rights—gained public attention during COVID-19 but predate it; see UK CMA (2018) and EU Transportation Committee (2016). The Spohr quote reflects Lufthansa's response to COVID-era demand reductions.

IV.C Resulting slot values and validation

While flights can involve multiple slots, (7) accounts all incremental profits to the departure slot at one airport—assuming matching arrival/departure slots are available at the same price for all carriers. Slot values are then scaled across 31 weeks, giving rise to a market for weekly slot pairs (WSPs).[31] Figure 3 plots the estimated WSP values and provides statistics. The mean (median) WSP value is €855k (€716k), with substantial variation, and the average contribution of wider network value measured with C_{ah} being 27 percent. Of the 861 slot values — covering all carrier-slot combinations with non-empty choice sets — 74% of the variation is across carriers within the same slot, reflecting different network opportunities. Air France exhibits notably higher values for some slots, indicating greater willingness to pay to retain its dominant Paris position. These higher values stem from comparative advantages associated with local network size: connectivity and frequency premia, cost efficiencies, and larger choice sets from more flights meeting slot-destination feasibility constraints.[32] An additional 20% of variation is across hours within the same day, reflecting both the slot-destination feasibility constraint and complementarities with existing flight banks during peak hours that enhance connectivity. Slots hold similar values at the two airports. The estimates imply prices of the same order of magnitude as prices paid for Heathrow slots.[33]

The welfare quantification that follows reflects the outcome of an allocatively efficient market-based mechanism for independent slot allocations, provided the model correctly identifies the highest-value firm for each slot. What matters is therefore the ranking of $\nu_a(s)$ across firms, not the level. Figure 4 plots the spread in slot values — the difference between the highest and second-highest values as a percentage of the second-highest. Despite substantial variation, a clear pattern emerges: coveted peak hours (6–8 am and 5–7 pm at CDG; 7–8 am and 8–10 pm at ORY) show much smaller spreads, indicating these times are consistently attractive to multiple carriers instead of facilitating profitable flights for individual carriers.

The large spreads overall — with a median of 133% — lend further credibility to the

[31] WSPs address within-airport pairing of arrival and departure rights and scaling across 31 consecutive weeks to cover a season. This packaging was applied in the never-adopted FAA JFK proposal. The potential bias from abstracting from complementarity between multiple slots on the same date at the same airport is small at the estimated model parameters and given the complementarity with carriers existing operations that the model does capture (Appendix E). The net bundle premium is at most 6% of the sum of individual slot values and it is of similar magnitude across carriers, whereas the spread between individual slot values of the winner and runner-up bidder is usually much larger as detailed below, leaving little room for bias.

[32] As Air France is not the preferred carrier in demand, this can be interpreted generally: existing network asymmetries carry through in the slot allocation unless the mechanism is designed to counteract them, for instance by favoring small bidders or newcomers.

[33] A perfect match is not expected given differences in airports, time periods, and non-random selection of transacted slots. However, slot values should be broadly comparable: Heathrow, Charles de Gaulle, and Orly are among the world's most slot-constrained airports, are geographically proximate, and share similar competitive environments with dominant carriers holding large slot shares. Appendix Figure B1 compares estimated Paris slot values with external Heathrow data.

Table 10: Observed vs. predicted network changes

	Main	No Cah	No hoard-ing	Drop random	Max share 0.6	Restricted
Precision (% predicted changes actually occurred)						
Expanded carrier-routes	50.0	50.0	55.6	57.1	69.2	100
Thinned carrier-routes -	70.0	70.0	70.0	40.0	68.4	78.9
Excess slot-constraint (predicted - actual % destinations)						
Expanded carrier-routes	-16.0	-13.1	-38.2	-9.1	17.4	10.7
Thinned carrier-routes	2.8	2.8	2.8	4.0	1.4	0.4

Notes. Precision: share of predicted carrier-routes that expanded or thinned that match actual S18–S19 changes in the global schedules data. *Excess slot-constraint*: difference between predicted and observed share of affected carrier-routes serving slot-constrained destinations.

welfare effects below: unmodeled factors affecting slot values differentially across carriers are unlikely to change the identity of the winning firm in most cases. However, in the bottom decile the spread is only 18%, where model misspecification is more likely to affect welfare estimates. The counterfactuals excluding a slot hoarding motive (setting $\rho_{a'}(s) = 0$ in (6)) yield larger spreads (*No Hoarding* in Figure 4), as easyJet and Lufthansa experience larger losses when losing a slot to a competitor. Without this effect, Air France would have a larger value advantage and win more auctions.

The estimates are also validated by comparing the predicted carrier-route changes to actual S18–S19 network changes (Table 10). I classify carrier-route combinations by whether more flights (per week) are scheduled in 2019 compared to 2018, and if so, I label the pair as "expanded" and otherwise as "thinned". Likewise, I classify carrier-routes based on whether I predict that flights are added or removed so that the two can be compared. The majority of predicted carrier-route expansions are observed (50–69% "precision"), and this number is (largely) unaffected by excluding broader network value or hoarding motives. Moreover, 68–70% of predicted carrier-routes with flight reductions are observed; well above the 40% achieved when shutting off the possibility for slot owners to optimize which flight to drop within the slot window. The "max share 0.6" column excludes carriers after winning 60% of auctioned slots, better matching observed schedule changes (69% precision), likely reflecting the first-come-first-serve waitlist mechanism limiting dominant carriers' expansion.[34]

Finally, Private values also permit construction of the Vickrey-Clarke-Groves (VCG)

[34]Together with countering the over-prediction of flight additions to slot-constrained destinations, this underscores that simulations are deliberately counterfactual—quantifying directional market changes when slots are reallocated based on values. For comparison, the final column of Table 10 restricts the counterfactual choice of $f_a(s)$ to flights on carrier-routes that actually expanded between S18 and S19, mechanically yielding 100% precision on additions; the modest improvement on contractions arises because correctly identified low-profit flights are no longer offset by simulated additions on the same routes. These insights are used to present welfare effects when restricting the choice of flight to the set of expanded carrier-routes in the schedules data. Appendix E provides further details on validation as well as complementarity among multiple departure slots.

allocation as a benchmark: with truthful reporting, the VCG-efficient assignment on a slot-by-slot basis awards each slot to

$$\arg\max_{a} \left[\nu_a(s) + \underbrace{\sum_{a' \neq a} \left(\mathbf{K}^*_{a'}\left(\mathcal{F}^{s(a)}\right) - \mathbf{K}^*_{a'}\left(\mathcal{F}^{s(\emptyset)}\right) \right)}_{\text{Externality of } a \text{ on competing bidders}} \right],$$

augmenting private slot values with the equilibrium profit changes imposed on competitors when a wins relative to the no-reallocation baseline $\mathcal{F}^{s(\emptyset)}$. The structural model recovers these externalities from the simulated profit matrix, permitting direct comparison of the VCG and private-value assignments. The two assignments agree in approximately 72% of slots. Divergences are systematic: the Air France group wins slots in 31 cases where VCG would prefer easyJet, and in 15 cases where VCG would prefer Lufthansa, so the VCG correction shifts some slot wins away from the dominant carrier.[35] Taken together, these results support that a private-value assignment based on $\nu_a(s)$ provides a close approximation to the socially efficient allocation, and I proceed with it as the basis for the welfare analysis.

V Welfare effects of market-based slot allocation

The market-based allocation is implemented by allocating each s to the carrier with the highest $\nu_a(s)$, which would be the result of an efficient auction mechanism or secondary market. After changing the flight network, adding the 245 chosen $f_a(s)$ of the winners and dropping the 245 $f^X(s)$, updating all network variables, and recomputing the equilibrium across all markets, welfare effects are quantified by comparing consumer surplus and firm profits against the baseline. As such, the welfare effects are based on changes in the network that go beyond the marginal changes on which the slot values are based. However, the results reflect directional effects when changing the slot allocation incrementally: 245 WSP make up about 6 percent of the sample ($245 \times 31/119,810 = 0.063$) and are valued in addition to the four carriers' existing operations. For simplicity, I refer to the allocation mechanism as an "auction", but the results are independent of specific mechanism design, implying also that the division of auction rents between bidders and the auctioneer via transaction prices is not identified.[36]

Air France initially owned 53% of the auctioned slots, easyJet 10%, Vueling Airlines

[35] The slot-by-slot VCG is a valid benchmark here because slot values are approximately additively separable conditional on (substantial) existing networks — the bundle premium from jointly winning slots at the same date and airport is at most 6% of the sum of individual values (Appendix E) — so the efficient joint assignment closely tracks the slot-by-slot one.

[36] Transaction prices resulting from the auction do not affect total welfare, as they are credited to the auctioneer (the airport or the slot coordinator), and, as a sunk cost at the flight level, do not affect fares.

Table 11: Slot allocation changes: unilateral direct effects

	All		ORY only		Restricted		Restr. no C_{ah}	
	WSP	ΔCS	WSP	ΔCS	WSP	ΔCS	WSP	ΔCS
Total (€M, markets with flight added & dropped)	245	290.4	119	147.7	245	131.6	245	109.4
By type (€/pp)	Drop	Add	Drop	Add	Drop	Add	Drop	Add
Business (5% pop.)	-0.87	66.81	-1.88	113.06	-1.88	12.88	-2.52	12.31
Leisure (95% pop.)	0.62	-1.31	0.47	-2.22	0.61	-0.55	0.59	-0.68
Reallocation patterns	% of total							
Different carrier wins	55.1	54.8	52.9	65.6	62.4	39.4	70.2	31.2
Competitor → AF	42.4	59.2	51.3	65.7	26.1	59.8	15.1	79.7
AF → Competitor	8.6	-2.8	0.8	0.1	21.6	21.0	38.8	51.5
Comp. → other comp.	4.1	-1.6	0.8	-0.2	18.0	-15.4	26.9	-56.3
Same carrier, diff. dest.	44.9	45.2	47.1	34.4	37.6	60.6	29.8	68.8
To constrained dest.	4.5	1.7	0.0	0.0	15.1	6.2	9.4	4.6
To unconstrained dest.	40.4	43.5	47.1	34.4	22.4	54.4	20.4	64.2
By winning carrier	% of total							
Air France	86.5	104.2	98.3	100.1	62.0	118.3	48.2	132.5
easyJet	7.3	-5.0	0.8	0.1	24.5	-14.3	34.3	-23.9
Lufthansa	5.7	0.9	0.0	0.0	0.4	0.3	0.8	0.4
Vueling Airlines	0.4	-0.1	0.8	-0.2	13.1	-4.2	16.7	-9.0
By time of day	% of total							
Morning rush (6–10am)	28.6	49.0	29.4	81.3	28.6	11.4	28.6	9.1
Mid day (11am–3pm)	28.6	20.7	29.4	0.4	28.6	47.3	28.6	53.1
Evening rush (5–8pm)	22.9	17.5	23.5	8.9	22.9	28.4	22.9	25.0
Late night (9pm–5am)	20.0	12.9	17.6	9.4	20.0	12.9	20.0	12.8
By destination change	% of total (241/245 change dest.)							
Short to short	12.2	-4.3	1.7	-0.1	58.0	-12.2	58.4	-26.9
Short to long	79.6	96.1	94.1	95.2	33.9	103.9	33.5	118.6
Long to short	1.2	0.1	0.0	0.0	3.7	-3.3	4.1	-6.8
Long to long	6.9	8.1	4.2	4.9	4.5	11.6	4.1	15.1
Flight characteristics	Drop	New	Drop	New	Drop	New	Drop	New
Load factor (%)	43.3	94.4	42.8	91.7	43.3	96.3	43.3	96.5
Baseline fare (€)	198	937	153	978	198	573	198	564
Capacity (seats)	168	442	175	450	168	322	168	318

Notes. Slot-level effects of reallocating each of the 245 WSPs to the highest-value carrier under four specifications: main ("All"); ORY slots only ("ORY only"); $f_a(s)$ restricted to routes that expanded S18-S19 ("Restricted"); same restriction excluding C_{ah} ("Restr. no C_{ah}"). Δ CS captures the change in consumer surplus in the two directly affected markets per slot; Table 12 reflects equilibrium including spillovers and interactions after all 245 slots are reallocated.

5%, Lufthansa 4%, and the remaining 49 carriers held the other third. Among the four carriers considered in the auction, Air France thus held 74%, easyJet 14%, Vueling Airlines 6%, and Lufthansa 5%. In the unconstrained scenario, Air France wins more auctioned slots than it loses, while Vueling Airlines loses most and easyJet approximately half of their auctioned slots (columns labelled "All" in Table 11). The associated consumer surplus

gains predominantly reflect slots being redirected from short-haul to long-haul routes — whether through carrier reallocation (55% of auctions) or incumbent reoptimization to a better-matched destination (45%) — and originates disproportionately in the morning rush hour, which accounts for 49% of consumer surplus gains despite representing 29% of slots. Looking only at ORY, where connecting traffic is minimal and hub-timing considerations are least relevant, the short-to-long pattern is even more pronounced (94.1% of reallocations) and morning slots account for 81.3% of consumer surplus gains.[37]

Under the restricted scenario, low-cost carrier expansion leads to some surplus-reducing flight changes for destinations only reachable from ORY (listed in Table B2), as reflected in the short-to-short destination change row. More generally, markets exclusively served from CDG—including San Francisco, Johannesburg, and Tel Aviv—show similar welfare effects as before; these are important expansion markets in both the restricted and unrestricted predictions, and indeed expanded in reality. Saint-Denis on Réunion is the most important unrestricted expansion market where no actual expansion occurred, as is the case for the other long-haul overseas territories served from ORY. Lufthansa's hands are tied the most in this scenario; only CDG-Munich flights can be chosen, and the fact that they don't win slots shows that these generate less profits than slot usage of competing carriers at CDG, while Munich is self-reported as a desirable expansion market for multiple carriers from ORY (Table A4).

Effectively excluding Lufthansa from the auction is necessary for Vueling Airlines to grow, offering more flights to Barcelona, Rome, and Sevilla. Likewise, the French slot coordinator awarded the ORY slots freed up by the Aigle Azur bankruptcy to Lufthansa, easyJet, and Air France but not to Vueling (see Tables A3-A4). However, Vueling had both capacity and ambition to expand its operations in Paris; it later received all 18 daily slot pairs at ORY that Air France was required to relinquish in 2021 as a remedy for competition distortions following state financial support, assigned by the European Commission via a beauty contest weighing total added seat capacity and connectivity potential of willing carriers with an operating base at ORY.

These results are robust to omitting the estimated outside network benefits (column "Restr. no C_{ah}"). Air France's hub preference determines whether its market share grows (from 53 to 62 percent) or shrinks (from 53 to 48 percent). When Air France is stripped of this advantage, low-cost carriers add flights that generate less consumer surplus on average than those they displace. Specifically, 58% of route substitutions swap one short-haul flight for another, and these substitutions produce a €27M net reduction in consumer surplus

[37] The efficiency gains therefore do not appear to be driven by omitted hub-timing effects but rather by genuine underutilization of short-haul slots at both airports. The drivers and magnitudes of CS gains appear similar, although Air France wins almost all slots at ORY and Lufthansa does not operate there.

across the 245 markets where flights are removed and added.

On the whole, only Air France expands in directions that improve consumer surplus– even under the most conservative specification. When the incumbent is predicted to switch only the destination, additional forces governing how flight networks are formed in reality — demand for return flights, fleet constraints, long-run demand projections, or finer time matching within the hour — might cause consumer gains to be overstated, but it remains remarkable that when especially easyJet and Vueling Airlines manage to be competitive, the resulting schedule changes are on average not consumer surplus-increasing.

Business travelers drive the aggregate gains: they gain €67 per person in markets receiving a new flight, falling to €13 in the restricted specification where fewer short-to-long morning reallocations occur. This fivefold difference confirms that business traveler gains are concentrated in precisely the slot type where the efficiency argument is strongest — morning slots redirected to long-haul routes. Leisure travelers, by contrast, are only modestly affected, losing at most €1.31 per person in markets with added flights and gaining at most €0.62 where one is removed, reflecting their price sensitivity and the small fare increases that accompany reallocation.

When including spillovers and interactions after all 245 slots are reallocated (for the restricted scenario), the slot auction generates €14.6M additional consumer surplus and €15.6M additional variable profits on average per market (Table 12). Only 167 out of 952 markets are directly affected by the 2×245 flight network changes, and all of the consumer surplus gain accrues there. Consumer surplus in pure spillover markets reduces by €0.8M on average. Without restricting flight choice to expanded routes, the spillover is positive (€0.7M, Table B5). This gap isolates cleanly the cost-side scheduling externality: since no flight is added or removed in spillover markets, it can only reflect larger winners operating a denser daily schedule, compressing costs and markups across their network. At 19% of a total gap in ΔCS of €7.8Mbetween the two scenario's, cost-driven network externalities make up a non-negligible share of the benefit of allocating more slots to larger carriers. The incremental approach leaves market structures largely unaffected, and, consequently, equilibrium price increases are modest (on average €0.46 across markets, see Table 12). In the unrestricted scenario, price increases are only slightly more pronounced, while still resulting in a positive surplus gap compared to the restricted simulations. This confirms that the majority of surplus gains comes from the dominant carrier having a larger network and the network expanding in higher-utility directions.

Another clear indication of efficiency gain is that dropped flights consistently have lower load factors (less than 50% vs more than 90%) and smaller aircraft than the added flights, both at baseline (Table 11). Airport runway capacity is used by substantially more passen-

Table 12: Aggregate welfare effects of slot changes

Averages across N markets	N	Δ CS (€ 10k)	ΔΠ (€ 10k)			Totals
All markets	952	14.6	15.6		ΔΠ	ΔC_{ah}
By market type				**Carrier profits**	(€ M)	(€ 100k)
Directly affected	167	86.9	98.0	Air France	116.0	6.8
Spillover	785	-0.8	-2.0	easyJet	21.8	3.0
Markets operated from				Lufthansa	-0.7	-0.6
CDG only	521	17.9	19.7	Vueling Airlines	11.0	1.5
ORY only	64	-23.2	17.5	**Aggregate effects**	Total/1000	Per flight
Both airports	367	16.5	9.3	Δ CO_2 (tons)	1.9	0.02
By route type				Δ Capacity (seats)	37.7	0.31
Short-haul (< 3,000 km)	536	10.4	3.3	Pollution (€)	1,682.3	26.00
Long-haul (≥ 3,000 km)	416	20.1	31.4	**Total welfare changes**		(€ M)
Price changes				a: ΔCS + ΔΠ		287.2
Δ Fare (Max/market €)	952	1.0		b: a − Pollution		272.5
Δ Fare (Avg/market €)	952	0.46		c: $b + \Delta C_{ah}$		290.0

Notes. Effects of reallocating 245 slots when limiting flights to carrier-routes that expanded between S18–S19, corresponding to the "Restricted" column in Table 11. Aggregate effects for the unrestricted scenario ("All") or when omitting C_{ah} are in Appendix Tables B5-B6. ΔCS stands for change in consumer surplus and ΔΠ for change in variable profits.

gers per slot, as after the 2021 beauty contest, and a total of 38k seats are added with the same capacity provided that aircraft is available. Larger aircraft comes with more pollution, which is computed as €250/Δtons CO_2 ×3 to capture the welfare costs of CO_2 and other global pollutants, increased by 21% to proxy for local air and noise pollution, matching the more detailed account of these environmental damages in Adler, Andreana and De Jong (2026) at baseline. This broad pollution measure lowers the welfare gains only by €1.68M to €272.5M.

About half of the welfare gains come from higher variable profits for the participating firms conform better slot utilization, with small losses for Lufthansa who is effectively the marginal competitor when restricting expansion opportunities.

VI Concluding remarks

This paper defines airport slot values as the incremental profits a carrier earns by operating a flight in a slot — estimated from a flight-level structural model that captures network spillovers in costs and demand, consumer preferences over departure timing, and the role of connecting traffic. Together with destination reachability and network value, these are the first-order sources of heterogeneity across slot values. Estimated values are validated against observed network changes between consecutive seasons.

The results show that reallocating a small share of congested-airport slots to existing carriers implies welfare gains driven by the addition of long-haul flights that disproportion-

ately benefit business travellers. Fare effects are negligible, confirming that gains stem from improved product offerings rather than changes in market power. While the resulting network changes use scarce airport capacity more efficiently, they come with increased pollution –although even when broadly valued, they amount to less than 1 percent of the welfare gains. Moreover, these gains are predominantly generated by the dominant carrier expanding on routes that increase consumer surplus rather than smaller carriers getting larger.

The central policy implication is clear: the current system of administrative slot allocation, sustained by use-it-or-lose-it hoarding, leaves significant value on the table. The effects of scaling up, slot packaging, and eligibility rules for new entrants are natural next steps for further study. The current estimates motivate that work by showing the welfare stakes of getting the design right are large.

References

Adler, Nicole, Gianmarco Andreana, and Gerben De Jong. 2026. "Estimating the societal value of airport slots." *Transportation Research Part B: Methodological*, 206: 103424.

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

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

Aryal, Gaurab, Charles Murry, and Jonathan W Williams. 2024. "Price discrimination in international airline markets." *The Review of Economic Studies*, 91: 641–689.

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

Bauer, Sebastian. 2025. "Competition and Welfare in Airport Slot Allocation." *Available at SSRN 5881882*.

Bergantino, Angela Stefania, Christian Bontemps, Mario Intini, and Ada Spiru. 2024. "Toward a consolidation of the European Airline Sector: the potential merger between ITA and LuftHansa."

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

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

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

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

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

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

Bontemps, Christian, Cristina Gualdani, and Kevin Remmy. 2023. "Price competition and endogenous product choice in networks: Evidence from the US airline industry." Toulouse School of Economics Working Paper 1415.

Bontemps, Christian, Gianmaria Martini, and Flavio Porta. 2026. "The effects of LCC subsidies on the tourism industry." *International Journal of Industrial Organization*, 104: 103242.

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

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

Borenstein, Severin, and Janet Netz. 1999. "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.

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

Brueckner, Jan K, Matthew E Kahn, and Jerry Nickelsburg. 2024. "How Do Airlines Cut Fuel Usage, Reducing Their Carbon Emissions?" *Economics of Transportation*, 38: 100358.

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

Cantillon, Estelle, and Martin Pesendorfer. 2006. "Combination bidding in multi-unit auctions." Working Paper, SSRN 958730.

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

Chen, Ertian, Lichao Chen, and Lars Nesheim. 2026. "Carbon regulation and competition in the European airline industry." Working Paper.

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

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

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

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

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

D'Haultfœuille, Xavier, Ao Wang, Philippe Février, and Lionel Wilner. 2022. "Estimating the Gains (and Losses) of Revenue Management." *arXiv preprint arXiv:2206.04424*.

Doraszelski, Ulrich, Katja Seim, Michael Sinkinson, and Peichun Wang. 2017. "Ownership Concentration and Strategic Supply reduction." National Bureau of Economic Research w23034.

Draganska, Michaela, Michael Mazzeo, and Katja Seim. 2009. "Beyond plain vanilla: Modeling joint product assortment and pricing decisions." *QME*, 7(2): 105–146.

Ershov, Daniel, and David J Salant. 2022. "Auctions and Mobile Market Competition: Evidence from European 4G Auctions." Working Paper, SSRN 4168166.

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

Fukui, Hideki. 2012. "Do Carriers Abuse the Slot System to Inhibit Airport Capacity Usage? Evidence from the US Experience." *Journal of Air Transport Management*, 24: 1–6.

Fu, Xiaowen, Tae Hoon Oum, Ruowei Chen, and Zheng Lei. 2015. "Dominant Carrier Performance and International Liberalization–The case of Northeast Asia." *Transport Policy*, 43: 61–75.

Gayle, Philip G. 2013. "On the Efficiency of Codeshare Contracts Between Airlines: Is Double Marginalization Eliminated?" *American Economic Journal: Microeconomics*, 5(4): 244–273.

Grigolon, Laura, and Frank Verboven. 2014. "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.

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

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

Hortaçsu, Ali, Olivia R Natan, Hayden Parsley, Timothy Schwieg, and Kevin R Williams. 2024. "Organizational structure and pricing: Evidence from a large us airline." *The Quarterly Journal of Economics*, 139(2): 1149–1199.

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

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

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

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

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

Li, Sophia, Joe Mazur, Yongjoon Park, James Roberts, Andrew Sweeting, and Jun Zhang. 2022. "Repositioning and market power after airline mergers." *The RAND Journal of Economics*, 53(1): 166–199.

Marra, Marleen. 2024. "A market for slots." CEPR Discussion Paper DP19442.

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

Milioti, Christina, and Amedeo R Odoni. 2024. "Airport "level" and market concentration in Europe." *Journal of the Air Transport Research Society*, 2: 100007.

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

Odoni, Amedeo R. 2020. "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.*

Park, Yongjoon. 2020. "Structural Remedies in Network Industries: An Assessment of Slot Divestitures in the American Airlines/US Airways Merger." Working Paper, SSRN 3710262.

Pouget, Lilian, Nuno Antunes Ribeiro, Amedeo R Odoni, and António Pais Antunes. 2023. "How do airlines react to slot displacements? Evidence from a major airport." *Journal of Air Transport Management*, 106: 102300.

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

Ribeiro, Nuno Antunes, Alexandre Jacquillat, António Pais Antunes, and Amedeo Odoni. 2019. "Improving slot allocation at Level 3 airports." *Transportation Research Part A: Policy and Practice*, 127: 32–54.

Ryan, Stephen P. 2012. "The costs of environmental regulation in a concentrated industry." *Econometrica*, 80(3): 1019–1061.

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

Suzuki, Yoshinori. 2006. "The net benefit of airline overbooking." *Transportation Research Part E: Logistics and Transportation Review*, 42(1): 1–19.

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

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

Yuan, Zhe, and Panle Jia Barwick. 2024. "Network Competition in the Airline Industry: An Empirical Framework." National Bureau of Economic Research w32893.

For Online Publication

A The global airport slot allocation policy

The International Air Transport Association (IATA) estimated in 2019 that 1.5 billion passengers departed from slot-constrained airports worldwide, constituting approximately 43% of all passengers in 2019.[38] As of March 2024, this number had increased to 50% of all passengers, indicating that demand has increased beyond pre-COVID-19 levels. These airports, also known as Level 3 airports, experience a structural imbalance between demand and capacity, where the need for airport facilities such as runway availability, arrival and departure assistance, and gate space consistently exceeds the available supply. As a result, it is imperative for all airlines and aircraft operators to obtain a slot allocated by a coordinator to access these airports.[39] Europe has the highest number of Level 3 airports, having in 2019 (2024) a total of 104 (114) out of the 204 (217) most congested airports during the summer season and 80 (84) out of 168 (187) during the winter season. In the US, the Federal Aviation Administration (FAA) managed slots for three Level 3 airports (JFK, LGA, and DCA) and coordinated slots for seven Level 2 airports (EWR, JFK, LAX, ORD, and SFO) in 2019; however, in 2024, only LGA was classified as Level 3. Most international connecting hubs are also classified as Level 3 airports.[40]

The allocation and monitoring of slots fall under the jurisdiction of national (or sometimes private) slot coordinators, who generally adhere to IATA's World Airport Slot Guidelines to standardize this process globally.[41] This scheduling process takes place twice a year to accommodate the industry's summer and winter schedules. In advance of each new season, regulators determine airport capacities for take-off and landing slots, as well as overall movements. Airlines submit slot requests to slot coordinators via a "slot clearance message" SSIM code.[42]

A crucial feature of the system is that airlines can keep their slots from the previous season if they utilize them at least 80 % of the time, which is known as the "use-it-or-lose-it" rule. Specifically, this concerns 80 % of the slot series, i.e., sequences of slots within a flight season. For example, a slot series would consist of all departure slots at 9 am on Monday during the 31 weeks of the summer season of 2018 at CDG airport. Any slot series for which at least 20% of the associated flights are canceled, along with new slots resulting

[38] Statistics in this paragraph are based on IATA publications; available at `https://www.iata.org/en/programs/ops-infra/slots/coordinated-airports/` and (from 2019) `https://www.iata.org/en/iata-repository/pressroom/fact-sheets/fact-sheet---airport-slots/`.

[39] IATA airport slot fact sheet.

[40] See Ribeiro et al. (2019) and FAA website.

[41] The latest WASG, including differences from previous guidelines, can be found here.

[42] For a beginner's guide to SSIM, see the European Airport Coordinators Association (EACA) document available here.

from capacity enhancements, or slots returned voluntarily, are placed into a "slot pool." The coordinators then allocate slots from this pool to fulfill requests from the waitlist.

Another crucial aspect is how pool slots are allocated to airlines on the waitlist. The World Airport Slot Guidelines (WASG) suggest allocating 50% of pool slots to new entrants.[43] Airlines are awarded airport-day status if, once their slot requests are accepted, they hold fewer than 7 slots at that airport on that day, for example, a maximum of 3 departures and 3 arrivals to/from any destination. Furthermore, within these categories, the guidelines prioritize i) requests that have been pending longer, ii) extending existing operations into year-round flights, iii) schedules that are effective for longer periods in the same season, iv) regular services over ad hoc services, and v) operations constrained by a (night) curfew at the other end of the route. However, even when adhering to all voluntary guidelines and local regulations where applicable, slot coordinators retain some discretion over waitlist requests. Most coordinators use one of approximately a dozen available slot coordination software programs.

In the period after the initial slot offering, slot coordinators and airlines communicate about slots via coded "SSIM" messages. These messages include codes for actions such as accepting the offered slot, requesting another slot nearby, remaining on the waitlist for a better slot, or indicating that no other slots are available. Additionally, the biannual global slot conference organized by the IATA plays a crucial role in the matching process. During this conference, airline representatives have the opportunity to meet face-to-face with slot coordinators. It is also established that airlines with new entrant status lose this status (for that airport/day) if they have not accepted a slot within one hour before or after the requested time by the end of the first day of the conference.

i. Slot constraints at CDG and ORY airports

In Europe, national slot coordinators legally operate within the rules set out by Council Regulation EEC 95/93 (updated in 2004 with EEC 793/2004). This regulation establishes common rules for the allocation of slots at European Community airports, very much in line with the IATA WASG guidelines.[44] COHOR, which is the slot regulator for France, is also

[43] The 50% rule is often cited as a fairness measure for new entrants. However, its effectiveness depends on the proportion of total slots that are available for reallocation through the pool. The guideline was slightly amended in the August 2019 updated WASG (previously called the World Slot Guidelines and available here). It now states that *"50% of the slots contained in the pool at initial slot allocation must be allocated to new entrant requests [...], unless new entrant requests are less than 50%. Similarly, 50% of the slots contained in the pool at the initial slot allocation must be allocated to non-new-entrant requests, unless such requests are less than 50%."* This update aims to distribute pool slots proportionally between new entrant requests and non-new entrant requests. Notably, increasing flight movements under new-entrant status may not always be advantageous, as this may limit scale economies associated with operating more flights from an airport on a given day.

[44] One difference is in the new entrant status; the European Slot Regulation, as amended in 2004, stipulates that an airline qualifies as a new entrant if it has fewer than 5 slots per day at the airport (two return flights)

Table A1: Known prices paid for traded Heathrow slots, 1998–2013 (CAPA, 2013)

Year	Buyer	Seller	Slots (daily, pairs)	Amount (Million £)	Daily slot value (Million £)
1998	BA	Air UK	4	15.6	3.9
2002	BA	BA Connect	5	13	2.6
2002	BA	SN Brussels	7	27.5	3.9
2003	BA	SWISS	8	22.5	2.8
2003	BA	United	2	12	6
2004	Virgin	Flybe	4	20	5
2004	Qantas	Flybe	2	20	10
2006	BA	BWIA	1	5	5
2007	BA	Malev	2	7	3.5
2007	BA	BA	7.3	30	4.1
2007	Virgin	Air Jamaica	1	5.1	5.1
2007	BMI	unknown	77.7	770	9.9
2007	unknown	Alitalia	3	67	22.3
2008	Continental	GB Airways, Alitalia, Air France	4	104.5	26.1
2013	Delta	unknown	2	30.8	15.4
2013	Etihad	Jet	3	46.2	15.4

Notes. Source: CAPA website (date of publication: 08-05-2013) and prices reported in GBP.

Table A2: Known prices paid for traded Heathrow slots, 2013–2017 (Odoni, 2021)

Year	Buyer	Seller	Slots (daily, pairs)	Amount (Million $)	Daily slot value (Million $)
2013	Etihad	Jet	3	70	23.3
2014	Etihad	Alitalia	5	80.46	16.1
2015	Turkish	SAS	1	22	22
2015	American	SAS	3	60	60
2015	Delta	Air France/KLM	6	276	46
2016	Oman Air	Air France/KLM	1	75	75
2017	Delta	Omni Air	75	37.5	

Notes. Source: Odoni (2020), based on Heathrow Airport holdings and prices reported in USD, with the exception of the price for the Etihad-Alitalia deal that is reported in euros (€60 million) and converted to dollars using the euro to U.S. dollar exchange rate for August 8 2014, i.e., 1.341. For comparability of trades reported in table A1, one can rely on the fact that both sources report the price paid for the Etihad-Jet slot trade in different currencies and use an implied exchange rate of 1 USD / 1.515 GBP.

c

subject to a French decree by the transport ministry that makes additional amendments to European regulation. The European air transport sector was liberalized in 1997, and the initial "historic" slot allocation was favorable to previously state-owned airlines. These so-called flag carriers are comparable to legacy carriers in the U.S. in terms of service, frequent-flyer programs, airport lounges, and tenure. Since 2000, the share of low-cost carriers has risen steadily to approximately 20% of all flights departing from France in 2013.[45]

The supply of slots at Level 3 airports Paris Charles de Gaulle (CDG) and Paris Orly (ORY) falls significantly short of meeting demand. To exemplify this issue at ORY, consider the 2019 bankruptcy of the Aigle Azur airline.[46] This exceptional circumstance led to COHOR reallocating a substantial pool of 9,868 slots at ORY midseason, equivalent to approximately 95 take-off and landing pairs for a weekly flight over a full year.[47] Despite the unusually large slot pool, only 18% of new entrant slot requests could be accommodated, and this was limited to just four out of the 19 entrant airlines. In essence, there was insufficient capacity for approximately 22,500 flights at ORY in 2020. Table A3 outlines all slot requests and awards by entrants for the affected period, along with all Aigle Azur slots granted to incumbent airlines, and Table A4 lists the proposed use of the slots for all entrant requests.

and intends to operate a flight to a route within the European Community that is operated by at most two other carriers.

[45] See page 22 of this OECD report on the basis of OAG data.

[46] The second airline bankruptcy in 2019, namely, XL Airways France, had a lesser impact. XL Air operated out of CDG but held only 1,488 slots in the 2019 summer season.

[47] Aigle Azur possessed 12,355 slots at ORY, yet the regulator DGAC reserved 20% of them for designated routes.

Table A3: Slot requests and awards at Orly, S20 season

		Entrant status[b]		
Airline	Requested[a]	*2b-i*	*2b-ii*	Awarded
Panel A: Not awarded — entrant status 2b-i only				
Air Caraïbes Atlantique	2,084	✓		
Air Nostrum	1,460	✓		
ASL Airlines France	2,190	✓		
Azur Brazilian Airlines	520	✓		
La Compagnie-Dreamjet	416	✓		
EasyJet[c]	730	✓		
Hop![d]	730	✓		
Ryanair	1,460	✓		
Siberia Airlines	730	✓		
Tassili Airlines	730	✓		
Ukraine International Airlines	730	✓		
Volotea	2,190	✓		
Wizz Air[e]	730	✓		
Chalair Aviation	520	✓		
EasyJet Switzerland	730	✓		
Vueling	2,602	✓		
Panel B: Not awarded — entrant status 2b-i and 2b-ii				
Anisec Luftfahrt GmbH	1,352	✓	✓	
Transavia Airlines C.V.[f]	624	✓	✓	
Panel C: Awarded — entrants				
EasyJet[c]	730	✓	✓	730
Hop![d]	2,920	✓	✓	1,460
Lufthansa	1,460	✓	✓	1,460
Wizz Air[e]	1,460	✓	✓	1,460
Total entrant slots	**27,722**			**5,110**
Panel D: Awarded — incumbents (requests unknown)				
Air Caraïbes International				1,460
Corsair International				1,460
Transavia France[f]				1,460
La Compagnie-Dreamjet				312
TAP Air Portugal[g]				66
Total incumbent slots				**4,758**

Notes. Source: COHOR website (published 05/12/2019). Requests were made for the S20 season but allocated based on the full year; airlines could begin operating immediately due to the unusual bankruptcy context. Only 18.43% of requested entrant slots were awarded to 4 of the 19 applicant airlines. The number of incumbent requests is unknown. a: Author's calculations based on proposed frequencies; a weekly flight corresponds to 104 slots and a daily flight to 730 slots. b: Entrant status under EU Slot Regulation 95/93 (amended 2004): 2b-i = fewer than 5 daily slots; 2b-ii = European route with fewer than three carriers. c: EasyJet appears twice because they requested slots under different entrant statuses. d: Hop! appears separately in Panels A and C for the same reason. e: Wizz Air appears separately in Panels A and C for the same reason. f: Transavia Airlines C.V. (Dutch low-cost subsidiary of Air France-KLM) requested 624 slots and was awarded none. Transavia France (French low-cost subsidiary of Air France-KLM) received 1,460 slots as an incumbent. g: TAP Air Portugal slots cover the summer season only.

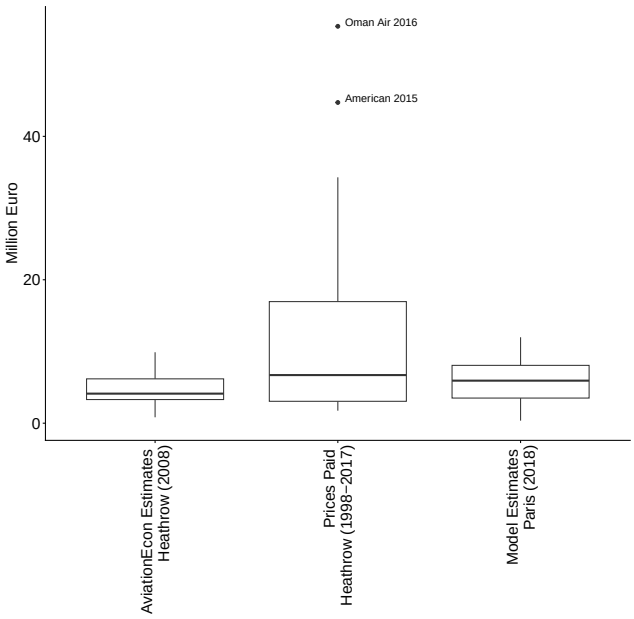
Table A4: Proposed and realized use of awarded slots, S20 season, Entrants only

Airline	Proposed use	Requested	Realized use
Not awarded			
Air Caraïbes Atlantique	3 weekly to Cayenne-Rochambeau; 1 daily to New York-Newark; 1 daily to São Paulo-Viracopos; 3 weekly to St-Maarten/Princess Juliana	2,084	—
Air Nostrum	2 daily to Madrid-Barajas	1,460	—
ASL Airlines France	2 daily to Algiers; 1 daily to Béjaïa, Belgrade, Djerba, Oran, Oujda, Tunis	2,190	—
Azur Brazilian Airlines	5 weekly to São Paulo-Viracopos	520	—
La Compagnie-Dreamjet	4 weekly to New York-Newark	416	—
EasyJet[c]	1 daily to London-Gatwick	730	—
Hop![d]	1 daily to Florence	730	—
Ryanair	2 daily to Dublin	1,460	—
Siberia Airlines	1 daily to Moscow-Domodedovo	730	—
Tassili Airlines	1 daily to Algiers	730	—
Ukraine Intl Airlines	1 daily to Kyiv-Boryspil	730	—
Volotea	2 daily to Algiers, Oran; 1 daily to Béjaïa, Constantine, Sétif, Tlemcen	2,190	—
Wizz Air[e]	1 daily to Bucharest, Vienna	730	—
Chalair Aviation	5 weekly to La Rochelle-Laleu	520	—
EasyJet Switzerland	1 daily to Hanover	730	—
Vueling	1 daily to Bergamo, Santiago de Compostela, Turin; 5 weekly to Bologna; 4 weekly to Göteborg-Landvetter, Pisa, Warsaw-Chopin, Zagreb; 3 weekly to Jerez; 6 weekly to Munich	2,602	—
Anisec Luftfahrt GmbH	13 weekly to Vienna	1,352	—
Transavia Airlines C.V.[f]	2 weekly to Bristol, Glasgow, Liverpool; 1 weekly to Gran Canaria; 3 weekly to Kraków, Santiago de Compostela	624	—
Awarded — entrants			
EasyJet[c]	1 daily to Glasgow	730	As proposed
Hop![d]	2 daily to Frankfurt; 2 daily to Munich	2,920	Frankfurt only
Lufthansa	2 daily to Munich	1,460	As proposed
Wizz Air[e]	1 daily to Budapest, Sofia, Warsaw-Chopin	1,460	As proposed

Notes. For details and c–f: see notes to Table A3. Proposals for incumbents slot usage are unavailable.

B Additional tables and figures

Figure B1: Comparison to 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 known as *Aviation Economics*, which is available upon request. Both their low-value and high-value estimates 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). The prices are also listed in Tables A1 and A2 in the Appendix. All prices are given in 2018 euros for a weekly slot pair for 31 weeks (as in S18). The model estimate of a slot price in Paris are based on estimated slot values $\nu_a(s)$ according to the main specification in (7). Prices are computed as 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, which is the average 2018 inflation rate in the European Union (Macrotrends)), taking only $\nu_a(s)$ for the winning bidder of each slot.

Table B1: Carriers by group

Air France	Flag Other	Non Flag Other
Air France	ANA	Aegean Airlines
KLM	Aer Lingus	Aigle Azur
Transavia France	Aeroflot	Air Austral
	Aeromexico	Air Caraibes
	Air Algerie	Air Transat
	Air Canada	Asiana Airlines
	Air China Limited	Atlasglobal
	Air Corsica	Corsair
	Air India	Delta Air Lines
	Air Madagascar	Jet Airways
	Air Mauritius	LATAMAirlines Brasil
	Air Serbia	United Airlines
	Air Seychelles	
	American Airlines	
	Austrian	
	British Airways	
	CSA	
	Cathay Pacific Airways	
	China Eastern Airlines	
	El Al	
	Emirates	
	Etihad Airways	
	Finnair	
	Iberia	
	JAL	
	Korean Air Lines	
	Middle East Airlines	
	Qatar Airways	
	Royal Air Maroc	
	SAS	
	SWISS	
	Singapore Airlines	
	TAP	
	TAROM	
	Thai Airways International	
	Tunisair	
	Turkish Airlines	
	Ukraine International	
	Vietnam Airlines	

Notes. The other carriers, Lufthansa, easyJet, and Vueling Airlines, are not grouped in estimation. *Non Flag Other* is the excluded category in estimation. *Flag Other* contains so-called "flag carriers" or (previously) state-owned national carriers.

Table B2: Destinations by airport availability and flight distance

CDG only		ORY only	Both	
Long-haul (≥3,000 km)				
Antananarivo	Mahe Island	Cayenne	Dzaoudzi	Philipsburg
Atlanta	Mexico City	Fort-de-France	Havana	Port Louis
Bangkok	Miami	Pointe-a-Pitre	Higuey	Saint-Denis
Beijing	Mumbai		New York	
Bengaluru	New York			
Bogota	Osaka			
Buenos Aires	Rio de Janeiro			
Cancun	San Francisco			
Caracas	Santiago			
Chicago	Sao Paulo			
Delhi	Seoul			
Hanoi	Shanghai			
Ho Chi Minh City	Singapore			
Hong Kong	Tokyo			
Johannesburg	Toronto			
Lima	Washington			
Los Angeles				
Short- and medium-haul (<3,000 km)				
Abu Dhabi	Istanbul	Agadir	Abidjan	Lisbon
Belfast	Krakow	Djerba	Ajaccio	Madrid
Belgrade	Kyiv	Fes	Algiers	Malaga
Boston	Lagos	Monastir	Amsterdam	Marrakech
Bristol	Libreville	Pisa	Athens	Marseille
Bucharest	Liverpool		Bamako	Milan
Cotonou	London		Barcelona	Montreal
Doha	Moscow		Bastia	Naples
Douala	Munich		Beirut	Nice
Dubai	Oslo		Berlin	Oran
Frankfurt	Saint Petersburg		Biarritz	Porto
Glasgow	Stockholm		Bordeaux	Prague
Hamburg	Warsaw		Budapest	Rabat
Helsinki	Zurich		Casablanca	Rome
			Copenhagen	Sevilla
			Dakar	Tel Aviv-Yafo
			Dublin	Toulouse
			Edinburgh	Tunis
			Faro	Venice
			Geneva	Vienna

Table B3: Estimated destination fixed effects

City	Airport	FE
Highest utility		
Lima	LIM	0.43
Dubai	DXB	0.46
Cotonou	COO	0.80
Dzaoudzi	DZA	0.80
Cancun	CUN	0.82
Port Louis	MRU	0.84
Higuey	PUJ	0.90
Libreville	LBV	1.16
Atlanta	ATL	1.30
Saint-Denis	RUN	2.43
Lowest utility		
London	LTN	-6.87
London	LHR	-6.23
Liverpool	LPL	-5.94
Marseille	MRS	-5.86
Bordeaux	BOD	-5.67
Milan	LIN	-5.64
Milan	MXP	-5.51
Belfast	BFS	-5.51
Barcelona	BCN	-5.51
Hamburg	HAM	-5.44

Notes. Destinations with highest and lowest fixed effects ("FE").

Table B4: Routes with changes in S18-S19 in schedule data

	Expanded		Thinned		Total
	N	Δ/wk	N	Δ/wk	Δ/wk
Total	558	4264.3	611	-1435.0	2829.3
By origin					
CDG	429	3735.3	515	-1263.2	2472.1
ORY	129	529.0	96	-171.7	357.3
By carrier					
Air France (CDG)	65	171.4	177	-241.8	-70.4
Air France (ORY)	42	98.8	41	-67.6	31.3
Lufthansa (CDG)	6	9.6	4	-1.5	8.1
Lufthansa (ORY)	0	0.0	0	0.0	0.0
Vueling (CDG)	4	16.9	6	-0.9	16.1
Vueling (ORY)	5	19.0	4	-3.3	15.8
easyJet (CDG)	23	14.9	14	-8.0	6.9
easyJet (ORY)	14	3.2	4	-3.6	-0.4
Other carriers					
CDG (N=82 carriers)	331	3522.4	314	-1011.1	2511.4
ORY (N=20 carriers)	68	408.0	47	-97.3	310.7

	Top expansions		Top thinning	
	Dest	Δ/wk	Dest	Δ/wk
1.	SVO	32.4	LIN	-14.7
2.	OPO	26.3	MAD	-11.2
3.	MXP	18.5	ALG	-10.5
4.	RAK	13.6	AMS	-10.2
5.	LIS	13.0	BCN	-8.3
6.	SVQ	8.4	ICN	-8.1
7.	SEA	7.7	IAD	-7.8
8.	NBO	7.6	FLR	-7.7
9.	SCQ	7.1	ATL	-7.4
10.	NCE	7.0	GVA	-6.2

Notes. A route is a carrier-origin-destination combination. A route that expanded (thinned) has more (fewer) flights/week scheduled by the operating carrier in the summer season of 2019 (S19) relative to 2018 (S18).

Table B5: Aggregate welfare effects of slot changes ("All" scenario)

Welfare effects, avg. across N markets	N	Δ CS (€ 10k)	ΔΠ (€ 10k)		Totals (€ M)	
All markets	952	22.4	28.2		ΔΠ	ΔC_{ah}
By market type				**Carrier profits**	(€ M)	(€ 100k)
Directly affected	158	131.3	179.8	Air France	258.5	13.9
Spillover	794	0.7	-1.9	easyJet	4.9	-0.5
Markets operated from				Lufthansa	6.5	0.4
CDG only	521	18.7	21.3	Vueling Airlines	-1.3	-0.7
ORY only	64	55.3	57.7	**Aggregate effects**	Total/1000	Per flight
Both airports	367	21.9	32.9	Δ CO_2 (tons)	3.4	0.03
By route type				Δ Capacity (seats)	67.1	0.56
Short-haul (< 3,000 km)	536	9.1	-2.1	Pollution (€)	3,113.8	25.99
Long-haul (≥ 3,000 km)	416	39.5	67.2	**Total welfare changes**		(€ M)
Price changes				*a*: ΔCS + ΔΠ		481.7
Δ Fare (Max/market €)	952	1.1		*b*: *a* − Pollution		478.6
Δ Fare (Avg/market €)	952	0.70		*c*: *b* + ΔC_{ah}		479.9

Notes. Effects of reallocating 245 slots when not limiting flights to carrier-routes that expanded between S18–S19, corresponding to the "All" column in Table 11. Δ CS stands for change in consumer surplus and ΔΠ for change in variable profits.

Table B6: Aggregate welfare effects of slot changes ("Rests. no C_{ah}" scenario)

Welfare effects, avg. across N markets	N	Δ CS (€ 10k)	ΔΠ (€ 10k)		Totals (€ M)	
All markets	952	12.7	16.1		ΔΠ	ΔC_{ah}
By market type				**Carrier profits**	(€ M)	(€ 100k)
Directly affected	171	74.1	99.0	Air France	104.7	0.0
Spillover	781	-0.7	-2.1	easyJet	35.5	0.0
Markets operated from				Lufthansa	-0.3	0.0
CDG only	521	18.0	19.0	Vueling Airlines	13.2	0.0
ORY only	64	-34.0	18.6	**Aggregate effects**	Total/1000	Per flight
Both airports	367	13.3	11.4	Δ CO_2 (tons)	1.8	0.02
By route type				Δ Capacity (seats)	36.7	0.31
Short-haul (< 3,000 km)	536	7.2	4.7	Pollution (€)	1,662.9	13.88
Long-haul (≥ 3,000 km)	416	19.9	30.8	**Total welfare changes**		(€ M)
Price changes				*a*: ΔCS + ΔΠ		274.1
Δ Fare (Max/market €)	952	1.2		*b*: *a* + SCC		272.4
Δ Fare (Avg/market €)	952	0.47				

Notes. Effects of reallocating 245 slots when limiting flights to carrier-routes that expanded between S18–S19 and omitting network opportunity costs C_{ah}, corresponding to the "Restr. no C_{ah}" column in Table 11. Δ CS stands for change in consumer surplus and ΔΠ for change in variable profits.

C Additional details on identification and estimation

The demand model is a random coefficient nested logit (RCNL) model (previously adopted in, e.g., Brenkers and Verboven (2006) and Miller and Weinberg (2017)) but with discrete consumer types (traveling for business or leisure), as seen 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), as described below:

$$\mathbb{E}[\xi(\theta^D)'Z^D] = 0. \tag{C1}$$

In this equation, $\xi(\theta^D)$ is the vector of flight-level structural error terms in consumer utility Equation (1) that equates 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 Equation (2), e.g., the implicit function that, in all markets, sets the following:

$$\hat{\mathbf{s}}_t = [S_{0t}(\mathcal{F}_t; \xi_t, \theta^D), \ldots, S_{Jt}(\mathcal{F}_t; \xi_t, \theta^D)]'. \tag{C2}$$

I apply the standard fixed-point algorithm of Berry, Levinsohn and Pakes (1995) to solve the implicit function, using the modified contraction mapping that Grigolon and Verboven (2014) showed 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.

The destination fixed effects are concentrated in the estimation algorithm, meaning that they are estimated for each set of candidate parameters $\tilde{\theta}^D$ for which the optimizer computes the GMM objective function. 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 as follows:

$$\hat{\theta}^D = \underset{\theta^D \in \Theta}{\arg\min} \frac{1}{N} g(\theta^D) A g(\theta^D)' \tag{C3}$$

$$g(\theta^D) = \tilde{\xi}(\tilde{\theta}^D)' Z^D. \tag{C4}$$

where $g(\theta^D)$ denotes the empirical moment condition (with destination fixed effects concen-

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

The solution to Equation (C3) is obtained with a constrained minimizer via 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 λ. Bounds are set only for the $\lambda \in [0.01, 0.99]$, $\gamma \in [0.01, 0.99]$, and $\alpha \in [-10, -0.001]$ parameters, and the results are always in the interior. The standard fixed-point algorithm outlined in Berry, Levinsohn and Pakes (1995) is adopted when solving for $\xi(\tilde{\theta}^D)$, with a tight convergence criterion of 1e-12. Overall, the performance of the described implementation of the estimator is surprisingly good given the scale of the problem. The results for the preferred specification are obtained in half an hour on a local machine.

Similar outcomes are obtained when first using the interior-point algorithm with a convergence criterion of 1e-5 and then switching to the sequential quadratic programming algorithm starting from the interior-point results. Both are large-scale algorithms suitable for the size of the problem. See Conlon and Gortmaker (2020) for further details about 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.

With the estimated $\hat{\theta}^D$ in hand, the marginal cost parameters $\hat{\theta^S}$ are obtained via instrumental variable regression. The marginal costs are obtained by subtracting the standard multi-product Bayes-Nash equilibrium profits of each flight given demand parameters ($B^*_{jt}(\mathcal{F}; \hat{\theta}^D)$) from prices and regressing the resulting costs $C_{jt}(\mathcal{F}_a)$ on $N^S_{jt}(\mathcal{F}_a)'\kappa^N$ and W_{jt}, instrumenting for endogenous quantities and iarcraft characteristics as described in Section II.E.

i. Instrument validity

Table C1 reports the instrument analysis underlying both the demand and supply estimates. Columns (1)–(2) present first-stage OLS regressions of price on the demand instruments and exogenous controls, with column (2) restricting the sample to raw (non-imputed) fares and adding destination fixed effects. Columns (3)–(5) are the analogous first stages for the supply system, regressing quantity, capacity, and wet-lease share on fleet composition and markup

instruments alongside the full set of exogenous controls.

The demand instrument set Z^D comprises fuel consumption and business-class share (cost shifters), the number of products in the market, network-based instruments (shares of monopoly and duopoly markets at the destination), BLP-style instruments based on flights' positions in characteristics space, and schedule-based instruments: the number of own and competing flights in the market (overall and from CDG), competing flights within a two-hour window, seats on competing flights, and average flight spacing at the destination.

A preliminary assessment using a plain logit specification (Berry, 1994) reveals that capacity and wet-lease instruments fail the Hansen J test under market-level clustering, consistent with airlines selecting aircraft on the basis of unobserved demand. These instruments are therefore excluded from the final demand instrument set. The schedule- and network-based instruments pass the overidentification test: the Hansen J statistic yields $p = 0.699$ in column (1) and $p = 1.000$ in column (2) (Frankfurt is excluded as a destination in column (2) since its instruments are not separately identified from destination fixed effects). This is consistent with the institutional setting: slot constraints at congested airports impose external restrictions on flight timing that are plausibly orthogonal to demand shocks, and the unit of observation is a single flight, so there is no mechanical correlation between flight frequency and product-level demand.

The Wu–Hausman test strongly rejects exogeneity of price in both specifications ($p = 0.000$), confirming the need for instrumentation. Weak instrument tests also confirm relevance ($p = 0.000$ in both columns), ruling out weak identification as a concern.

The supply instrument set comprises fleet composition variables: average capacity and the share of small aircraft and wet-leased aircraft on the carrier's other routes, and the analogous averages for competing carriers in the market (columns (3)–(5)). These instruments are strongly relevant—the Sanderson–Windmeijer weak instrument test rejects weakness for capacity and wet-lease share separately ($p = 0.000$ in both cases)—and pass the joint Hansen J overidentification test ($p = 0.973$, with standard errors clustered at the route level and quantity as the dependent variable). The Wu–Hausman test rejects exogeneity of the supply-side endogenous variables jointly ($p = 0.000$), confirming that quantity, capacity, and wet-lease share must be instrumented for in the supply system. Note that for these initial supply-side tests, I take Quantity rather than mark-up as the dependent variable, the latter being computable with the structural demand estimates only. Hansen J statistics are also reported alongside the structural estimates.

Table C1: Analysis of demand and supply instruments

"First stage" (OLS regression)	Demand-side		Supply-side		
Dependent variable:	Price		Quantity	Capacity	WetLease
	(1)	(2)	(3)	(4)	(5)
Fuel Consumption	0.248***	0.357***	-0.264	-3.026***	-0.004***
	(0.009)	(0.015)	(0.341)	(0.157)	(0.001)
Business Class percentage	-0.003*	0.001	1.628***	1.416***	-0.003***
	(0.001)	(0.002)	(0.042)	(0.028)	(0.000)
Monopoly Markets (at Destination)	0.577***	0.260*	14.926***	5.371***	-0.008
	(0.047)	(0.107)	(0.990)	(1.080)	(0.007)
Duopoly Markets (at Destination)	0.127***	-0.073	0.290	0.631	0.090***
	(0.029)	(0.062)	(0.689)	(0.692)	(0.004)
Number of products in Market	0.000*	0.001***	-0.016***	0.008***	-0.000***
	(0.000)	(0.000)	(0.005)	(0.001)	(0.000)
Nr. Destinations	-0.088***	-0.087*	-3.779***	-13.177***	0.155***
	(0.017)	(0.037)	(0.703)	(0.538)	(0.003)
Number own flights from CDG in Market	-0.001***	-0.003***	0.067***	0.023***	0.000*
	(0.000)	(0.001)	(0.009)	(0.004)	(0.000)
Number competing flights from CDG in Market	0.002***	0.004***	-0.061***	-0.019***	0.001***
	(0.000)	(0.000)	(0.005)	(0.003)	(0.000)
Number competing flights in Market	0.020***	-0.014+	-1.588***	-0.832***	0.001*
	(0.003)	(0.007)	(0.116)	(0.088)	(0.001)
Nr. own flights in 2 hour window in Market	-0.013+	-0.052**	-6.738***	-3.272***	-0.001
	(0.008)	(0.018)	(0.229)	(0.176)	(0.001)
Nr. competing flights in 2 hour window in Market	-0.009	-0.001	0.889***	-0.822***	0.011***
	(0.007)	(0.014)	(0.257)	(0.207)	(0.001)
Fleet: Avg capacity on carrier's other flights			-0.038***	0.042***	0.001***
			(0.005)	(0.004)	(0.000)
Fleet: Share of carrier's other flights small aircarft			-65.061***	-31.895***	0.250***
			(1.434)	(0.993)	(0.006)
Fleet: Share of carrier's other flights wet leased			109.160***	41.049***	1.018***
			(3.743)	(2.223)	(0.013)
Fleet: Avg capacity on competing flights in Market			0.251***	0.930***	-0.001***
			(0.015)	(0.004)	(0.000)
Fleet: Share of competing flights in Market small aircraft			-16.683***	6.834***	-0.478***
			(2.352)	(0.780)	(0.005)
Num.Obs.	119,810	48,277	119,794	119,794	119,794
R2 Adj.	0.911	0.888	0.687	0.785	0.261
Destination FE	✓	✓	✓		
Raw fares only		✓			
"Second stage" diagnostics					
Dependent variable:	$ln(s_j) - ln(s_0)$			Quantity	
Weak instruments (p)	0.000	0.000		0.000	0.000
Wu-Hausman (p)	0.000	0.000		0.000	
Hansen J (p, clustered)	0.699	1.000		0.973	

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

Notes. Columns (1)–(2) report the first-stage regressions for the demand IV estimator, showing the relationship between endogenous prices and the demand instruments conditional on exogenous demand variables, with or without removing imputed fares. Column (3) is an OLS first stage for the supply system, regressing quantity on fleet composition and markup instruments alongside the full set of exogenous controls. Columns (4)–(5) are analogous first stages for capacity and wet-lease share, respectively. The Weak instruments (p) values report the Sanderson–Windmeijer test for Capacity and WetLease separately. Wu-Hausman and Hansen J statistics are reported jointly for the supply system (cols 3–5): the Wu-Hausman test evaluates the exogeneity of the supply-side endogenous variables, and the Hansen J test assesses overidentification of the instrument set, with standard errors clustered at the market level.

D Additional details about the role of time and connectivity preferences

Consumers' time preferences are a crucial factor for differentiation in slot values. I analyze different specifications to understand how estimated time preferences interact with scheduling and connectivity of flights (Table D1). Specifications with departure or arrival hour fixed effects reveal a pattern that peaks in the morning and dips roughly 12 hours later. I fit a cosine functional form to capture such preferences while smoothing out non-monotonic hourly effects:

$$\texttt{Time Cosine} = \cos\big(\frac{\pi}{12}(\text{Hour} - h)\big), \tag{D1}$$

with h the most preferred hour. The main specifications peak at $h = 10$am for departure time and $h = 7$am for arrival time, fitting the estimated hourly dummies.

Specifications with departure time preferences yield negative coefficients on departure connectivity despite producing economically sensible Lerner indices of 0.45–0.47, closely aligned with Air France's reported gross margin of 0.48.[48] This sign reversal reflects a mechanical feature: departure connectivity counts feeder arrivals 1–6 hours prior to departure, but for a preferred 10am departure, this window (4am–9am) captures few arrivals due to night-time restrictions.

By contrast, arrival hour specifications preserve positive and significant connectivity coefficients. Arrival connectivity counts onward departures 1–6 hours after arrival; a 7am arrival accesses the 8am–1pm window with abundant traffic. When introducing type-specific arrival coefficients, the business segment expands to 16 percent of the population, now characterized by stronger preferences for early morning arrivals.

The choice of connectivity measure introduces further trade-offs. Simple connectivity, counting all flights within the temporal window regardless of routing efficiency, yields uniformly positive coefficients, unlike the constrained measure requiring connections not exceed 1.5 times direct flight time. The constrained measure may be overly restrictive, excluding viable connections that passengers value.

The diversion analysis highlights an undesirable consequence of restricting passenger types to have identical time preferences. Under such restrictions, business and leisure travelers substitute in roughly equal proportion to other morning flights. Models with type-specific preferences obtain differentiated diversion: business types divert 2.9 times as much as leisure types to morning flights. This makes both the size and timing of slot holdings relevant for market power.

To illuminate how passenger composition and network structure shape demand patterns, I estimate the main specification excluding CDG flights (column 4 of Table 4). The share of

[48] Source: Macrotrends.net.

business travelers more than doubles at ORY ($\hat{\gamma}_1 = 0.125$ versus 0.052), consistent with the airport preference coefficients in other specifications where business consumers prefer ORY (negative coefficient on CDG) while leisure consumers prefer CDG. Time-sensitive business travelers likely value ORY's proximity to central Paris and faster security screening, whereas leisure travelers prefer CDG's larger choice set and superior amenities.

The larger arrival time coefficient for both passenger types at ORY suggests that timing is more relevant at this airport. ORY operates a predominantly short-haul, point-to-point network serving domestic and European destinations; early morning arrivals are particularly valuable, enabling business travelers to reach their destination for a full working day. At CDG, the hub structure with waves of connecting banks dilutes these pure schedule preferences.

The connectivity coefficients underscore this distinction. Connectivity at the departure airport drops sharply at ORY (0.108 versus 0.344). CDG passengers—including those connecting through the Air France hub—place substantial value on departure airport connectivity. With a different mix of legacy and low-cost carriers at ORY (Table B1), Air France becomes the preferred carrier from this airport.

Table D1: Alternative time preference specifications

	Time Preferences					Economics			Diversion $\frac{8am \to [6\text{-}10am]}{8am \to [>10am]}$			Connections	
Model	Cosine	Fixed Effects	Arrival	2-type Coef.	All Connections	Lerner Index	Bus. Share	% Inelastic	Business	Leisure	Ratio	Departure Coef.	Arrival Coef.
Main (Arr cosine 7am, 2type, simple connec)	✓		✓	✓	✓	0.42	0.05	0.1	1.46	0.50	2.9	+	+
Dep FE		✓				0.47	0.15	0.0	0.65	0.55	1.2	−	+
Dep cosine (10am)						0.45	0.05	0.2	0.71	0.57	1.2	0	+
– 2type	✓			✓		0.56	0.11	1.3	1.31	0.28	4.6	−	+
– 8am	✓					0.45	0.05	0.1	0.71	0.57	1.2	0	+
– 8am + 2type	✓			✓		0.79	0.15	17.3	1.72	0.07	25.8	−	+
Arr FE		✓	✓			0.45	0.04	0.4	0.70	0.58	1.2	+	+
Arr cosine (7am)	✓		✓			0.47	0.04	0.2	0.70	0.57	1.2	0	+
– 2type	✓		✓	✓		0.53	0.16	0.6	1.13	0.29	4.0	−	+
– simple connec.	✓		✓		✓	0.39	0.19	0.0	0.67	0.52	1.3	+	+
– 8am + 2type	✓		✓	✓		0.52	0.14	0.7	1.61	0.26	6.2	−	−
No time						0.39	0.15	0.0	0.65	0.55	1.2		

Notes. Comparing models with different time preference specifications against the *Main* specification. *Dep FE* = departure hour dummies, *Arr FE* = arrival hour dummies, *Cosine* specifications based on equation (D1). Diversion statistics compute predicted diversion from 8am flights to 6–10am flights as a ratio of diversion to flights after 10am. Final columns indicate significantly positive (+), negative (−), or insignificant (0) willingness to pay for connectivity.

E Validation of the Slot Value Model

i. Predicted vs. Observed Network Changes

To validate the assumptions underlying the slot value model, I compare predicted carrier-route changes to the changes carriers actually made between the 2018 and 2019 observed flight schedules (Table 10). The reported *precision* takes all carrier-routes where flights are predicted to be added (or removed) and measures the share of those that correspond to actual network changes: carrier-routes with higher (or lower) weekly flight frequencies in S19 compared to S18, accounting for the one additional week in S19 and selecting only carrier-routes that existed in both periods. The *excess slot-constraint* measure captures whether predicted changes disproportionately involve slot-constrained destinations relative to actual changes.

Precision for carrier-routes that expanded ranges from 50–69%, while carrier-routes that thinned achieve 68–70% except under random dropping (40%), in which case the operator cannot optimize which flight to expire within a slot targeted for reallocation. The main specification underestimates network expansion to slot-constrained destinations by 16 percentage points. This gap narrows under the scenario where slot rights expire randomly without substitution to lower-profit flights ("drop random"). The difference arises because the main specification systematically targets lower-profit flights for removal, and these are disproportionately at slot-constrained destinations. Under random expiration, slots at constrained destinations are released in proportion to their frequency share, undoing less of any subsequent shift toward these airports. This pattern is consistent with the inefficient slot use driven by the use-it-or-lose-it rule: carriers retain relatively low-profit flights to slot-constrained destinations to preserve slot rights, making these flights unlikely targets under profit-based removal but equally likely under random expiration. That random dropping performs substantially worse at predicting which carrier-routes are actually shed (40% vs. 70% precision) confirms that carriers respond to flight-level profitability as estimated rather than dropping routes arbitrarily.

The "max share 0.6" column excludes carriers after winning 60% of slots up for auction while maintaining the incremental character of the model. Better matching of observed schedule changes (69% precision) likely reflects the first-come-first-serve waitlist mechanism that limits dominant carriers' expansion in practice. These results underscore that the simulations are deliberately counterfactual, quantifying directional market changes when historic rights and slot constraints are removed, rather than predicting equilibrium network configurations under the current system.

For comparison, the final column restricts the counterfactual to flights on carrier-routes

that actually expanded between S18 and S19, mechanically yielding 100% precision on additions. The modest improvement in predicting route contractions arises because correctly identified low-profit flights are no longer offset by simulated additions on the same carrier-routes. The main body of the paper uses this information to present the welfare effects achievable under system-wide slot constraints, where new routes require available slots at both origin and destination.

Taken together, these exercises confirm that the slot value model provides reasonable predictions despite its simplicity: additions follow a highest-profit heuristic subject to slot feasibility, removals target lowest-profit flights within slot windows, and C_{ah} accounts for network motives not captured by direct-route profitability alone. A formal product choice model with flight-specific fixed costs could improve accuracy, though estimation with individual flight products and slot allocation constraints remains a challenge for future research.

ii. Complementarity across slot packages

One potential concern is that complementarities across banks of connecting flights might be missed since the mechanism allocates slots individually. Efficient market-clearing mechanisms can break down in the presence of complementarities (e.g., Milgrom, 2000; Cantillon and Pesendorfer, 2006). Three types are already incorporated in the Weekly Slot Package item definition. A fourth — complementarity across multiple reallocated WSPs, which may vary with carriers' existing slot holdings — is not, and its magnitude needs to be assessed empirically.

Indeed, the presented simulations assign slots sequentially, with each carrier bidding its marginal value $\nu_a(s)$ for slot s. This value is defined in (7) as the difference in total profits for the carrier with and without the slot, incremental to the carrier's network at baseline ($\mathcal{F}_a$). Writing $v(s+\mathcal{F}_a)-v(\mathcal{F}_a)$ for this quantity, and omitting the carrier subscript on $v(\cdot)$ to focus on within-carrier structure, the mechanism values a *bundle* of slots $\{s_1, s_2\}$ for a single carrier as the sum of individual bids:

$$\big[v(s_1+\mathcal{F}_a)-v(\mathcal{F}_a)\big]+\big[v(s_2+\mathcal{F}_a)-v(\mathcal{F}_a)\big]. \tag{E1}$$

Because both bids are computed against the same baseline $\mathcal{F}_a$, this sum need not equal the true bundle value $v(s_1+s_2+\mathcal{F}_a)-v(\mathcal{F}_a)$. The difference is exactly the bundle premium (or "cross-partial"):

$$\text{Bundle premium } \Delta(s_1,s_2)=v(s_1+s_2+\mathcal{F}_a)-v(s_1+\mathcal{F}_a)-v(s_2+\mathcal{F}_a)+v(\mathcal{F}_a). \tag{E2}$$

When $\Delta > 0$ (supermodular) the implemented mechanism *undervalues* the bundle of slots s_1 and s_2 by Δ; when $\Delta < 0$ (submodular) it *overvalues* the bundle. This misevaluation can cause misallocation only if it is large enough to change the ranking of bids across carriers. It is therefore relevant to establish whether $|\Delta|$ is small relative to the marginal value of a slot for all carriers.

Analytical bounds on the bundle premium

Computing $\Delta(s_1, s_2)$ exactly requires solving the full pricing equilibrium under four network configurations: $\mathcal{F}_a$, $\mathcal{F}_a + s_1$, $\mathcal{F}_a + s_2$, and $\mathcal{F}_a + s_1 + s_2$. In each configuration, equilibrium quantities and prices adjust to reflect the changed network variables (*FlightstoDest*, *Spacing*, *NrFlights*), and profits depend on all of these jointly. Instead, I characterise $\Delta(s_1, s_2)$ analytically under a first-order approximation: I hold equilibrium quantities fixed at their baseline values when computing how adding one or two slots changes profits through each model component. Under this approximation, the bundle premium decomposes additively into three terms corresponding to the three network variables in the model. Each term has a determinate sign that follows from the functional form alone, independent of the approximation. The magnitudes of the three components contributing to the bundle premium, and the net amount (R_{net}), are approximations that omit the second-order effect of quantity adjustment. Concretely: adding a slot raises *FlightstoDest*, which raises demand and equilibrium quantities for all flights on that route, which in turn changes the cost contributions of *Spacing* and *NrFlights* through the changed passenger volumes. The analytical ratios below treat these cross-effects as zero.

Throughout, let $\hat{\beta}_{\log f} = 0.071$ denote the estimated demand coefficient on $\log(\textit{FlightstoDest})$, and let $\hat{\kappa}_{\text{spa}} = 0.067$ and $\hat{\kappa}_{\text{fli}} = -0.238$ denote the estimated marginal cost coefficients on *Spacing* and *NrFlights*/100 respectively. Let n_d denote the actual number of flights the carrier operates to destination d on the relevant date ($n_d = 10 \times \textit{FlightstoDest}$) and N the actual total number of departures from the origin airport ($N = 100 \times \textit{NrFlights}$). See tables 4 and 6 in the paper for the coefficients, taken from the main specifications.

Demand term ($\log$ *FlightstoDest* in utility). When s_1 and s_2 both serve destination d on the same date, adding them jointly raises n_d by two instead of one twice from the same base. Holding equilibrium quantities fixed, the bundle premium of the demand contribution to profits is proportional to the second difference of $\log n_d$:

$$\Delta_{\text{dem}} = \hat{\beta}_{\log f} \cdot \log\left(\frac{n_d(n_d+2)}{(n_d+1)^2}\right) \leq 0, \tag{E3}$$

which always non-positive regardless of the approximation: the denominator $(n_d + 1)^2$ ex-

ceeds the numerator $n_d(n_d+2)$ by exactly 1 for any n_d, so the logarithm is strictly negative.[49] The demand term therefore contributes *submodularity*: the log-frequency gain from a second slot on the same route, $\log[(n_d+2)/(n_d+1)]$, is smaller than that from the first, $\log[(n_d+1)/n_d]$, because log is concave. The ratio of the second to first difference,

$$\frac{\Delta_{\text{dem}}}{\log[(n_d+1)/n_d]} = \frac{\log[n_d(n_d+2)/(n_d+1)^2]}{\log[(n_d+1)/n_d]}, \tag{E4}$$

is exact and depends only on n_d; $\hat{\beta}_{\log f}$ cancels. Note that this is the ratio of the second to first log-frequency difference, not a ratio of profit increments: translating it into a profit ratio requires the full equilibrium, which is what the first-order approximation bypasses.

For different-destination pairs, or same-destination pairs on different dates, $\Delta_{\text{dem}} = 0$ exactly: $\log(\textit{FlightstoDest}_d)$ and $\log(\textit{FlightstoDest}_{d'})$ are distinct arguments in the utility function, so adding a slot to d' leaves the frequency term for d unchanged.

Spacing term (*Spacing* in marginal cost). *Spacing* is the average of adjacent hourly departure gaps on a given date, mathetmatically identical to $R/(N-1)$, where $R = t_N - t_1$ is the operating range (hours from first to last departure on the date) and N is the total flight count. Adding a new slot affects *Spacing* differently depending on whether it falls within the existing window or extends it.

Let $R_0 = S_0(N-1)$ denote the baseline operating range with S_0 the baseline Spacing and N the baseline flight count. Let R_1, R_2, R_{12} denote the operating ranges after adding s_1 alone, s_2 alone, and both together, with $\delta_k = R_k - R_0 \geq 0$ the corresponding range extensions ($\delta_k = 0$ if slot k is interior to the existing window). The resulting spacings are $S_1 = R_1/N$, $S_2 = R_2/N$, $S_{12} = R_{12}/(N+1)$.

Under the first-order approximation, the profit contribution of the spacing term for M flights is $-\hat{\kappa}_{\text{spa}} \cdot S \cdot \sum_f q_f$, where $\hat{\kappa}_{\text{spa}} = 0.067$ (in 100 euros per passenger per hour of spacing) and $\sum_f q_f$ is the total number of equilibrium passengers across all flights. Holding passengers fixed at their baseline values and approximating $\sum_f q_f \approx N\bar{q}$ where $\bar{q}$ is the average equilibrium passenger count per flight, the bundle premium of the spacing contribution is:

$$\Delta_{\text{spa}} = -\hat{\kappa}_{\text{spa}}\bar{q}\big[S_{12}(N+2) - S_1(N+1) - S_2(N+1) + S_0 N\big]. \tag{E5}$$

Substituting $R_k = R_0 + \delta_k$ and separating the baseline from the range-extension components:

$$\Delta_{\text{spa}} = \underbrace{-\frac{2\,\hat{\kappa}_{\text{spa}}\,S_0\,\bar{q}}{N(N+1)}}_{\text{interior compression, } \leq 0} - \hat{\kappa}_{\text{spa}}\bar{q}\underbrace{\left[\frac{\delta_{12}(N+2)}{N+1} - \frac{(N+1)(\delta_1+\delta_2)}{N}\right]}_{\text{range-extension correction, } \geq 0}. \tag{E6}$$

[49] Expanding: $(n_d+1)^2 - n_d(n_d+2) = 1$. This is the AM-GM inequality applied to $a = n_d$ and $b = n_d+2$: their arithmetic mean n_d+1 strictly exceeds their geometric mean $\sqrt{n_d(n_d+2)}$ whenever $a \neq b$.

Table E1: Model implications for slot complementarity

	Network variables			Bundle premium			
	$\log(\tilde{f}_d)$	$\tilde{N}$	$\bar{S}$ (h)	R_{dem}	R_{spa}	R_{fli}	R_{net}
At baseline networks							
Air France	-1.005	1.765	0.113	-16.0%	*-1.1%*	*+0.6%*	*+0.6%*
Lufthansa	-0.445	0.127	1.179	-12.4%	-14.6%	+7.6%	+5.6%
Vueling Airlines	-1.401	0.130	1.547	-20.7%	-14.3%	+7.4%	+5.0%
easyJet	-1.495	0.311	0.553	-23.6%	-6.2%	+3.2%	+3.1%
At post-reallocation networks							
Air France	-1.005	1.763	0.113	-16.0%	*-1.1%*	*+0.6%*	*+0.6%*
Lufthansa	-0.447	0.126	1.183	-12.4%	-14.7%	+7.6%	+5.6%
Vueling Airlines	-1.394	0.131	1.540	-20.7%	-14.2%	+7.4%	+5.0%
easyJet	-1.494	0.311	0.551	-23.6%	-6.2%	+3.2%	+3.1%
In polar case: each carrier hypothetically wins all 245 slots							
Air France (wins all)	3.239	4.215	0.047	*-0.4%*	*-0.5%*	*+0.2%*	*+0.2%*
Lufthansa (wins all)	3.229	2.577	0.054	*-0.4%*	*-0.8%*	*+0.4%*	*+0.4%*
Vueling Airlines (wins all)	3.226	2.580	0.072	*-0.4%*	*-0.8%*	*+0.4%*	*+0.4%*
easyJet (wins all)	3.220	2.761	0.060	*-0.4%*	*-0.7%*	*+0.4%*	*+0.4%*

Notes. First-order bundle premium $R_{\text{net}} \equiv \Delta(s_1, s_2)/\left[\nu_a(s_1) + \nu_a(s_2)\right]$ and its decomposition into three channels, evaluated at observed carrier-level network statistics under three scenarios. $\log(\tilde{f}_d)$ is the log number of flights by the carrier to a given destination on a given date (in tens); $\tilde{N}$ is the total number of flights by the carrier to any destination on a given date (in hundreds); $\bar{S}$ is the average spacing in hours between adjacent departures of the carrier on a given date. Columns (A) report statistics at baseline networks; columns (B) after the implemented reallocation under the Restricted scenario (see Section V); columns (C) in the polar case where each carrier in turn hypothetically wins all 245 auctioned slots, giving an upper bound on R_{net} throughout the auction. R_{dem}, R_{spa}, and R_{fli} give the contribution of $\log(\tilde{f}_d)$, *NrFlights*, and *Spacing* respectively to $\Delta(s_1, s_2)$, each as a share of $\nu_a(s_1) + \nu_a(s_2)$; see equations (E3)–(E6). All ratios are computed under a first-order approximation that holds equilibrium quantities fixed at baseline values; the spacing term additionally assumes interior slot insertion and will be updated once slot-time data are incorporated. Entries below 2% in absolute value are in italics.

The general expression (E6) decomposes the spacing bundle premium into two components. The interior compression term is negative: each slot that falls within the existing window $[t_1, t_N]$ splits one gap and reduces *Spacing*, with the second slot compressing it by less than the first because the schedule is already denser — submodularity. The range-extension correction is non-negative: slots that extend the operating window reduce *Spacing* by less than interior slots, pushing the spacing bundle premium toward zero or positive. The net sign and magnitude therefore depend on the distribution of auctioned slot times relative to the existing departure window, which I assess directly in the simulations reported in Table E1 using observed values of S_0, δ_1, δ_2, and δ_{12}.

Net bundle premium. The demand term contributes submodularity; the *NrFlights* term contributes supermodularity; the sign of the *Spacing* term depends on the distribution of auctioned slot times relative to the existing departure window and cannot be determined analytically. The net $\Delta(s_1, s_2)$ from (E2) therefore cannot be signed analytically without further information. Table E1 evaluates $\Delta(s_1, s_2)$ and its ratio to the sum of stand-alone slot

values, $R_{\text{net}} \equiv \Delta(s_1, s_2) \,/\, \big[v(s_1+\mathcal{F}_a)-v(\mathcal{F}_a)+v(s_2+\mathcal{F}_a)-v(\mathcal{F}_a)\big]$, at observed carrier-level network statistics under three scenarios: baseline networks (column A), post-reallocation networks (column B), and a polar upper bound in which each carrier hypothetically wins all 245 auctioned slots (column C).

At baseline, R_{net} is positive for all four carriers — +0.6% for Air France and +3–6% for the three sparse carriers — so the Nr. Flights complementarity dominates in all cases and the mechanism undervalues joint bundles by these margins. Two features of the results bound the scope for this to affect allocation outcomes.

The first is that the bundle premium is small for the carrier whose bids are decisive. At Air France's baseline network of $N \approx 177$ daily flights, the $1/N$ decay in both cost ratios leaves a net premium of only +0.6%. The sparse carriers face higher premia of +3–6%, driven by their large spacing ($\bar{S} \approx 0.6$–1.5 hours) and small N (12–31 flights), but their stand-alone bid levels are often substantially lower than Air France's as the latter wins most auctions and the median spread is 133%, as detailed in the main body of the paper. The polar case confirms that supermodularity does not accumulate as a carrier acquires more slots: when any single carrier wins all 245, both cost ratios decay to near zero and all premia fall below 0.5%.

The second is that since all carriers face positive bundle premia of similar order of magnitude, correcting for supermodularity moves all bids in the same direction and is unlikely to change the allocation ranking. Misallocation requires the corrected ranking $(1+\rho_A)\,\nu_A$ vs. $(1+\rho_B)\,\nu_B$ to differ from the uncorrected ranking ν_A vs. ν_B, which demands that the difference in premia $|\rho_A - \rho_B|$ times the stand-alone value exceeds the winner's margin of victory. Given that all premia lie in the range +0.6% to +6%, the maximum differential is around 5 percentage points, and this would need to exceed a margin of victory that is typically large relative to the slot value.

F Summary of simulation approach

This section summarizes the counterfactual simulations. I begin by simulating baseline equilibrium profits for the four airlines across all flights under the observed schedule, incorporating a cost of overbooking to curb this at maximum 15%. A slot s is a departure-airport–hour–day-of-week combination in which there are flights at baseline, yielding 245 unique values.

Outer loop: For each slot s, starting from the baseline data, I draw one flight at random from among those departing in s — on average 488 flights per slot across 55 airlines. The drawn flight's airline determines which of its flights is displaced from s: in the main specification, the flight with the lowest baseline profit among all its departures in s; in a robustness check, the drawn flight itself (I refer to this as "random dropping"). The displaced flight is recorded as $f^X(s)$.

Inner loop. For each of the four airlines, I determine the set of potential flight choices accounting for slot-destination feasibility and prior network changes, and select the feasible flight with the highest baseline profit as its preferred addition to slot s, denoted $f_a(s)$. I then update all network variables affected by the removal of $f^X(s)$ or the addition of $f_a(s)$, and recompute counterfactual equilibrium profits for each airline. *End inner loop.*

For each airline, I retain the profit difference between adding its preferred flight to the schedule and the worst-case outcome for that airline among each of its three competitors doing the same; in one robustness check I ignore the profit effects of what competitors would do altogether (I refer to this as "no hoarding"). In the main specification, I subtract the estimated network opportunity cost $\hat{C}_{ah}$ to obtain the net slot value $v_a(s)$ according to (7); one robustness check sets these costs to zero. Slot s is assigned to the airline with the highest $v_a(s)$, and its selected flight is recorded as $f^W(s)$. *End outer loop.*

Post-reallocation. Starting from the baseline schedule, I drop the 245 flights in f^X, add the 245 flights in f^W, update all relevant network variables, recompute the full equilibrium under the reallocated network, and calculate the implied welfare effects.