Transportation Research Part B 206 (2026) 103424

Contents lists available at ScienceDirect

Transportation Research Part B

journal homepage: www.elsevier.com/locate/trb

Estimating the societal value of airport slots

Nicole Adler [a], Gianmarco Andreana [b,a,*], Gerben De Jong [c,d]

[a] Hebrew University Business School, The Hebrew University of Jerusalem, Mount Scopus, Jerusalem, 91905, Israel
[b] University of Bergamo, Department of Economics, Via dei Caniana 2, Bergamo, 24127, Italy
[c] SEO Amsterdam Economics, Roetersstraat 29, Amsterdam, 1018 WB, Netherlands
[d] Vrije Universiteit Amsterdam, School of Business and Economics, De Boelelaan 1105, Amsterdam, 1081 HV, Netherlands

A R T I C L E I N F O

Keywords:
Game theory
Discrete choice
Slot value
Aviation networks

A B S T R A C T

This research integrates discrete choice modelling and game theory to estimate the societal value of airport slot capacity. We first estimate heterogeneous passenger preferences for fares, frequencies and flight directness and then embed these estimates into a game-theoretic framework of competing airlines. Our framework captures the propagation effects of slot regulation across networks, the distributional impacts on passengers and airlines and environmental externalities from global emissions and local pollution. We apply this framework to assess the implications of tightening slot capacity at hub airports in North America and Europe, regions that collectively account for around 50% of worldwide revenue passenger kilometres in 2024. Social welfare declines when slot capacity is reduced because diminished connectivity and higher fares erode consumer surplus. Airline profitability is only marginally affected under mild slot reductions because increased market power raises per-passenger revenues, which offset lower passenger volumes. The environmental gains from slot reductions do not outweigh the loss in consumer surplus. However, from the perspective of a local regulator, slot reductions may increase social welfare if global environmental impacts are internalized and the social cost of carbon is sufficiently high, particularly at hubs with a large share of long-haul routes. Our findings highlight the importance of accounting for hub location and network effects in the cost-benefit analysis of slot capacity management.

1. Introduction

The setting and allocation of airport slot capacity is a key issue in aviation policy, particularly as the industry faces growing demand, limited infrastructure and increasing environmental pressure. Slots, which grant permission to take off or land at a specific time, are central to managing connectivity and capacity at congested airports, shaping market access and competition at more than two hundred slot-constrained airports worldwide (Belobaba et al., 2015).[1] Independent slot coordinators seek to balance operational efficiency, equitable market access and environmental concerns, yet the coordination of slot capacity is complex and current practices are often opaque.[2] As a scarce resource, slots allow airlines to secure market entry, establish strategic routes and maintain international networks. Their private economic value to airlines is evident from Oman Air's purchase of a pair of slots at London Heathrow in 2016

* Corresponding author.
E-mail addresses: nicole.adler@mail.huji.ac.il (N. Adler), gianmarco.andreana@unibg.it (G. Andreana), g.dejong@seo.nl (G. De Jong).

[1] In summer 2025, there were 215 slot coordinated airports worldwide, according to IATA, accessed on 12.6.2025: https://www.iata.org/en/iata-repository/pressroom/fact-sheets/fact-sheet-airport-slots/

[2] Currently, slots are allocated through the IATA Worldwide Airport Slot Guidelines (Iata, 2023). The guidelines are developed and updated by the IATA Slot Committee, which includes representatives of airlines, airports and regulatory authorities. The guidelines establish a globally recognized framework for managing airport slots. One key aspect is the 'use-it-or-lose-it' rule, which typically requires utilizing a slot at least 80% of the time

https://doi.org/10.1016/j.trb.2026.103424
Received 16 December 2024; Received in revised form 8 December 2025; Accepted 6 February 2026
Available online 16 February 2026

for $75 million (Aviation Week Network, 2016). However, their limited availability, combined with oligopolistic market structures and environmental externalities, complicates efforts to assess their broader societal value.

Against this backdrop, recent policy discussions in Europe have focused on tightening slot capacity to reduce environmental impacts such as greenhouse gas emissions and noise (Reuters, 2023). Concurrently, academic research has long highlighted distortions created by existing capacity management systems, such as barriers to fair competition (Brueckner, 2009; Ciliberto and Williams, 2010; Fukui, 2019; Verhoef, 2010). Without a clear understanding of the marginal societal value of slots, it is difficult to assess whether current allocation mechanisms and policies serve the interests of the broader community, including passengers and the aviation supply chain. We fill this gap by developing a methodological approach to estimate the societal value of airport slots and apply it to assess the impacts of slot reductions at major hub airports. By jointly modelling passenger preferences, airline competition and environmental externalities, we quantify the costs and benefits of alternative policies, including local priorities such as connectivity and sustainability objectives. The framework yields welfare-consistent marginal slot values and decomposes impacts on consumer surplus, producer profits and global and local pollutants, enabling an evidence-based assessment of slot-restriction policies at major hubs.

Assessing whether slot capacity restrictions serve societal interests requires understanding how passengers value connectivity, the strategic airline responses and resultant environmental externalities. Passengers' choices and willingness-to-pay for air travel are central determinants of the welfare impacts of aviation (Berry et al., 2006; Berry and Jia, 2010) and are shaped by factors such as routing, schedules and fares, which vary substantially across passenger groups and regions. The strategic behaviour of airlines shapes competitive market dynamics. Dominant carriers may restrict frequencies to raise fares above competitive levels and leverage incumbent positions to block competitors and deter entry, which reduces consumer welfare and undermines efficient capacity allocation. Equally important is the relatively large and hard-to-abate environmental impact of aviation, which has become a central policy concern. While optimal slot constraints have been extensively studied as a tool to address airport delays, focusing on the trade-off between congestion externalities and market power, we examine a slot capping policy as an instrument for addressing noise and emissions, taking into account both global pollution and local effects on air quality and the health of inhabitants living near airports. Finally, auction-based approaches have been proposed to derive slot values and improve transparency (Ball et al., 2020; Rassenti et al., 1982), but most existing work scales poorly and centres on private values to airlines, omitting broader welfare components and externalities. We address these gaps by developing a network-based framework that integrates passenger demand, competitive airline behaviour, connectivity and environmental externalities into a unified cost-benefit analysis.

From an applied theoretical perspective, we bridge the gap between empirical observations and theoretical models in order to estimate the impact of airport slot constraints on aviation markets. The primary methodological contribution lies in developing a cost-benefit framework that incorporates heterogeneous demand estimates into a theoretical framework that explicitly models airline strategic behaviour across a network. This offers decision-makers a foundation for evaluating the potential impacts of slot constraints and other policy interventions in an empirically grounded, reasonably large-scale network model. From an empirical perspective, we estimate demand coefficients separately for economy and business passengers across the European, North American and Transatlantic markets using six different nested logit models. These models capture geographical differences between the key drivers of air travel demand, including fares, flight frequency and route directness. We find that price sensitivity differs across regions and segments-North American economy passengers are more price sensitive than Europeans, while the opposite holds in business travel-and that Europeans strongly prefer more direct services. From a policy perspective, we present a case study applying our approach to examine the marginal value of airport slots at major European and North American hubs, focusing on their implications for passenger surplus, airline profitability, the environment and overall social welfare. We find that airlines often benefit from mild slot reductions, as the resulting increase in market power more than offsets the loss of traffic. However, the savings in global and local emissions and noise pollution are insufficient to cover the consumer welfare loss from higher fares and lower frequencies. Our results also show that slot values and the welfare effects of slot-restriction policies are a function of the location of the hub where restrictions are applied.

The remainder of this paper is structured as follows. Section 2 discusses the existing literature. Section 3 provides the empirical model and the resulting demand parameters for the different markets and passenger types. We introduce our game-theoretic framework in Section 4. Section 5 presents an application of our game to a network covering two continents. Concluding remarks and future directions are provided in Section 6.

2. Literature search

We have organized the review to cover the theoretical and empirical literature examining game-theoretic models of airline competition, discrete-choice models analysing passenger demand, policies addressing the environmental impacts of aviation, and the economics of airport slot allocation.

2.1. Game-theoretic models of airline competition

Hansen (1990) develops a non-cooperative model of airline competition in a hub-and-spoke network, where airlines compete on flight frequency while fares remain fixed, showing that such a model is able to identify quasi-equilibrium states that closely mirror

in order to keep the slot in the subsequent scheduling period. This approach prioritizes continuity and stability for incumbent carriers but tends to limit access to new entrants.

real market conditions. Hong and Harker (1992) introduce a computable air traffic network equilibrium model to estimate the value of landing slots for airlines. When landing slots are endogenized, most slots are allocated to the airline with the highest benefit, concentrating market share and profits. Adler (2005) formulates a two-stage competitive model in which network configurations that maximize airline profitability are chosen in the first stage and frequencies, fares and aircraft size in the second stage. The findings suggest that competitive equilibria in airline networks are challenging to sustain due to network economies and capacity constraints, including slot allocation systems.

Adler et al. (2010) construct a dynamic game where airlines compete among themselves and against a high-speed rail operator in terms of fares and service frequency, solving the resulting non-linear maximization problem in the European Union context to evaluate the overall impact of high-speed rail infrastructure investments on social welfare. Hansen and Liu (2015) design two models capable of predicting competition between two symmetric airlines that differ only in their frequency competition structure. Employing a nested logit specification, they demonstrate that the results from their analytical framework are able to replicate equilibrium outcomes. More recently, Wang et al. (2022) design an equilibrium model to study airline frequency competition at slot constrained airports, taking into account demand flows and airline alliances, demonstrating that a Nash equilibrium may not always exist in pure strategies but does arise when the setting permits mixed strategies. For a comprehensive review of transport market modelling using game-theoretic principles, see Adler et al. (2021).

A large empirical and theoretical literature examines how alliances, antitrust immunity and joint ventures affect fares, service quality and welfare in international aviation. A consistent finding is that cooperation on interline itineraries lowers fares by mitigating double marginalisation: allied or immunised carriers jointly set fares for connecting trips rather than independently marking up their respective segments, leading to discounts of roughly 15–25% relative to non-aligned pairs (Brueckner, 2003; Brueckner and Whalen, 2000). Brueckner and Singer (2019) show that joint venture arrangements typically deepen this integration through metal-neutrality, whereby partners pool revenues and become indifferent to which airline operates the flight, further aligning incentives and often yielding pricing behaviour indistinguishable from that of a single carrier. By contrast, cooperation on overlapping gateway-to-gateway routes can weaken competition, though measured fare effects range from negligible to moderate depending on market size and regulatory carve-outs (Brueckner, 2009). Beyond pricing, Adler and Hanany (2016) theoretically show that coordination may also influence frequencies because alliances increase producer surplus through joint scheduling, while limited code-sharing may maximise consumer surplus when passengers value frequency and schedule convenience.

Relative to this literature, our framework differs in two key respects. First, rather than estimating reduced-form fare effects route-by-route, we rely on structurally estimated passenger demand, which allows us to recover heterogeneous willingness-to-pay for routes, frequencies and fares across a network. Second, we estimate a multi-airline Nash equilibrium in both fares and frequencies across all markets simultaneously, capturing substitution patterns, network spillovers and strategic interactions that route-level regressions or partial-equilibrium models cannot accommodate.

2.2. Discrete-choice models of passenger demand

Understanding passenger behaviour is necessary to evaluate the welfare effects of aviation policies. Berry et al. (2006) develop a framework with heterogeneous products, diverse consumer preferences and economies of density, providing a foundation for structural analysis of airline markets. Using a structural differentiated-products model, Berry and Jia (2010) analyse the US airline industry and show that shifts in demand, changes in marginal costs and the emergence of low-cost carriers account for much of the decline in legacy carrier profits. Similarly, Gayle (2013) develop a structural econometric model to quantify the effects of double marginalization and code-sharing, while Barnhart et al. (2014) estimate passenger flows and delays using a multinomial logit model. Together, these studies provide empirical foundations for modelling passenger behaviour and demand characteristics.

A large empirical literature has estimated air travel demand elasticities and documented systematic differences across market segments and regions. Early interpretative surveys (e.g. Oum et al., 1992) emphasise that business travellers tend to be less price-sensitive than leisure travellers, and that elasticities vary with the availability of substitute modes. More recent work uses booking or route-level data to quantify these patterns. Granados et al. (2012) show that price elasticities differ across booking channels and between business and leisure passengers, Mumbower et al. (2014) demonstrate how elasticities vary across advance-purchase horizons and Morlotti et al. (2017) document substantial variation in elasticities across European low-cost routes by destination type, advance purchase, and seasonality. Complementary gravity-based studies such as Boonekamp et al. (2018) relate route-level passenger volumes to GDP, population, tourism, and low-cost carrier presence across European airport pairs. However, most of this work either focuses on a single region, specific carrier type or market segments, making cross-region comparisons indirect and sensitive to modelling choices. Our contribution is to estimate a structural, product-level discrete-choice system with a common specification for six combinations (economy and business passengers in the European, North American and transatlantic markets), which yields directly comparable elasticities for fares, frequency and route directness.

Most empirical work on air travel demand originates from the United States, largely due to data availability. European studies remain comparatively scarce. Birolini et al. (2020) develop an integrated origin-based demand model for air transportation, while Birolini et al. (2021) address airline network planning with demand-supply interactions. Most recently, Birolini et al. (2023) extended this integrated approach to airport slot allocation, demonstrating that prioritizing passenger-centric metrics, such as minimizing missed connections, yields significantly different allocation outcomes than traditional flight-centric optimization. However, no study has yet examined the structural demand differences between the US and European markets within a consistent empirical framework, a gap we address in this paper

Discrete-choice studies also show that passengers value schedule convenience and service quality in addition to fares. Higher frequencies reduce schedule delay, the difference between preferred and scheduled departure time, while fewer connections, shorter detours, and better on-time performance improve travel time and reliability (Cho et al., 2017; Hansen, 1990; Hess et al., 2007). Business travellers generally exhibit a higher value of time and are more sensitive to schedule delay, frequency and directness than leisure passengers (Cho et al., 2017; Proussaloglou and Koppelman, 1999). At the network level, airlines often use frequency as a competitive instrument, leading to relatively high frequency served by small aircraft. This behaviour reflects both passenger preference for schedule convenience and competitive dynamics, with implications for congestion, capacity use and environmental impacts (Givoni and Rietveld, 2009). In line with this literature, our demand specification includes fares, the logarithm of flight frequencies to capture diminishing marginal utility of additional flights, a non-stop indicator and detour distance as key product characteristics.

2.3. Environmental impacts of aviation

Policies addressing the environmental impacts of aviation are receiving increasing attention. Brueckner and Zhang (2010) analyse the effects of airline emission charges on airfares, service quality and aircraft design. Fageda and Teixidó (2022) provide evidence on the impact of the EU Emissions Trading Scheme on the aviation sector, which has helped mitigate emission growth but not reduce the absolute values. De and G (2022) examines how emission pricing has contributed to earlier capital replacement decisions in the relevant aviation markets. A complete treatment of aviation's environmental impact should take into account not only global emissions but also local emissions affecting noise and air quality which in turn impact the health of inhabitants living in close proximity to an airport (Adler et al., 2013; Fageda and Flores-Fillol, 2025). In addition, non-CO_2 effects such as contrails imply that radiative forcing from aviation exceeds that of CO_2 alone, motivating the use of a multiplier to approximate CO_2-equivalent impacts (Lee et al., 2021).

2.4. Economics of airport slot allocation

The airport slots literature offers insights into the trade-offs involved in different allocation methods. Borenstein (1988) highlights a significant disparity between private and social incentives in the slot allocation process, establishing a foundation for subsequent welfare analyses. Verhoef (2010) argues that regulation should address two key market failures; the failure to internalize congestion and the high fares arising from market power. He further posits that tradeable slots may enhance efficiency if congestion externalities are substantial, but may be less effective if market power is the more significant concern.

Several studies examine how slot constraints affect airline behaviour and welfare. Vaze and Barnhart (2012) develop a game-theoretic model to examine how competing carriers respond to frequency constraints at a single airport, finding that a small reduction in slots leads to a large reduction in frequencies, reduced delays and increased airline profits. However, their analysis does not account for network effects or overall social welfare. Swaroop et al. (2012) show that wider implementation of slot controls at US slot-constrained airports would improve traveller welfare, finding that current slot caps at US airports are insufficiently restrictive, though successful implementation requires addressing practical challenges such as efficient slot allocations and assessment of the impact on local communities. Ciliberto and Williams (2010) show that control over airport gates at large hubs leads to higher hub premiums and ticket prices. Similarly, Fukui (2019) leverage the removal of slot constraints at Newark to demonstrate that slot controls reduce competition and increase airfares. Silva and Verhoef (2013) demonstrate that addressing congestion externalities through per-flight tolls and correcting for market power inefficiencies via per-passenger subsidies may achieve first-best outcomes, offering insights relevant to slot pricing and capacity management. Adler and Yazhemsky (2018) evaluate the marginal impact of changes in slot availability at congested airports from the perspectives of airports, airlines and passengers, finding that relaxing slot caps in Europe could increase movements at the cost of marginally higher delays, thereby improving overall social welfare, whereas reducing slots at the four most constrained US airports would further limit infrastructure utilization, leading to diminished overall social welfare.

Slot auctions have been widely suggested as a first-best solution to raise transparency and welfare of the slot-allocation process. As Rassenti et al. (1982) demonstrate, auction mechanisms could ensure that airlines reveal their true valuation of airport slots whilst enhancing overall welfare, thus achieving optimal outcomes. Despite the appealing characteristics of auctions, their implementation is far from trivial, particularly in a network context. Ball et al. (2020) develop quantity-contingent auctions for airport slot allocation, while Bichler et al. (2023) design airport time slot auctions complying with the IATA scheduling guidelines. However, these models primarily focus on airline valuations, often neglecting broader welfare considerations including externalities. Moreover, the inherent complexity of the combinatorial optimization problem in allocating slots across multiple airports limits these models, which are primarily designed for and tested on small-scale scenarios. We refer the reader to Zografos et al. (2017) for a comprehensive review of the literature on airport capacity and optimal slot allocation.

Collectively, our contribution connects four strands of the literature. First, we embed empirically estimated, heterogeneous passenger demand in a game-theoretic network model of airline competition under slot constraints. Second, we link this to a social welfare function that explicitly includes both global emissions and local air-quality impacts. Third, we evaluate slot capacity policies in a multi-airport network, in order to analyse how changes at a major hub propagate across markets. Finally, we provide welfare-based slot estimates that extend existing auction-based approaches, which are typically airport specific and focused solely on private airline values.

3. Empirical strategy

We model demand using a nested logit discrete choice model in the spirit of Berry (1994). Beyond computational convenience, the nested logit model yields a closed-form solution for aggregate market shares, which enhances the applicability of our estimates in subsequent theoretical work. We estimate the model separately for economy and business passengers, given their differential preferences over frequency and fares. The specification accounts for the outside option of not flying, correlation among unobservable characteristics of the flight alternatives (inside options) and the endogeneity of the fare variable. To address the latter, we use well-known instruments, including the number of firms or products in the market (e.g., Berry et al., 1995) and a Hausman-style cost shifter exploiting the panel structure of our data (Hausman et al., 1994). The estimation uses 2019 Official Airline Guide (OAG) data on passenger traffic, airfares and schedules. To address widespread missing fare data (30 to 50% of the fare data is missing depending on the market), we supplement the OAG data with the Airline Origin and Destination Survey (DB1B) from the US Department of Transportation and apply product-market averaging and imputation methods detailed in Appendix A.

3.1. Utility and market share

Consumers θ are assumed to be utility maximizing agents that select airline product a in the travel market from departure city i to arrival city j. The consumer utility function is specified as:

$$U_{ij\theta a} = Z_{ij}\alpha + X_{ija}\beta + \xi_{ija} + \epsilon_{ij\theta a}, \tag{1}$$

where Z_{ij} denotes a vector of exogenous market characteristics which explain the utility of flying relative to the outside good; X_{ija} is a vector of observed (potentially endogenous) product characteristics that explain the choice between products as well as the utility relative to the outside good; ξ_{ija} represents consumer-averaged unobserved product quality; and $\epsilon_{ij\theta a}$ is the distribution of consumer preferences around this average.

To take into account the correlation among the unobservable characteristics of the flight alternatives, we model market shares according to a nested formulation. We group products into $\mathcal{B}$ groups, $b \in \{0, 1\}$, where $b = 0$ represents the outside good (not flying), while $b = 1$ is associated to the flight alternatives (inside goods).[3] Products are assigned to groups by $b(a)$ and the set of products in group b is denoted as $\mathcal{J}_b$. This nested specification is as follows:

$$U_{ij\theta a} = Z_{ij}\alpha + X_{ija}\beta + \xi_{ija} + \bar{\epsilon}_{ij\theta b(a)} + (1-\rho)\bar{\epsilon}_{ij\theta a}, \tag{2}$$

where the consumer preferences around the average utility, $\epsilon_{ij\theta a} = \bar{\epsilon}_{ij\theta b(a)} + (1-\rho)\bar{\epsilon}_{ij\theta a}$, follows the typical nested logit formulation i.e., as $\rho \to 0$ the nested formulation collapses to the multinomial logit, whereas $\rho \to 1$ implies perfect within group correlation.

The aggregate market shares are estimated as the product of two logits, the within-group logit and the across-group logit:

$$\hat{m}_{ija} = \frac{\exp\left(\frac{V_{ija}}{1-\rho}\right)}{V_{ijb}} \frac{V_{ijb}^{(1-\rho)}}{1 + V_{ijb}^{(1-\rho)}} \tag{3}$$

where $V_{ija} = Z_a\alpha + X_{ija}\beta + \xi_{ija}$ is the consumer-averaged utility for product a in market ij; and $V_{ijb} = \sum_{a\in\mathcal{J}_{ijb}} \exp\left(\frac{V_{ija}}{1-\rho}\right)$.

Taking logs of the flight alternatives market shares and differencing with the log of the outside good market share, leads to a linear estimable equation (Berry, 1994):

$$\ln(m_{ija}) - \ln(m_{ij0}) = Z_{ij}\alpha + X_{ija}\beta + \rho\ln(m_{a|ijb(a)}) + \xi_{ija}, \tag{4}$$

where $m_{a|ijb(a)}$ is the (endogenous) within group share of product a.

We specify the exogenous and endogenous market and product characteristics, i.e. the variables in the vectors Z_{ij} and X_{ija}, as follows: the market characteristics include a constant, great circle distance between the origin and endpoint cities and gross domestic product per capita in the origin and endpoint cities. The product characteristics are ticket fare, frequency, directness and airline specific intercepts. We take the logarithm of frequency to reflect the decreasing returns of additional frequencies, as is common in the literature (Hansen, 1990). The abovementioned characteristics represent the typical flight product attributes used in air transport demand research.

Eq. (4) contains two endogenous variables. The first is the within-group market share, $m_{a|ijb(a)}$. Naturally, a higher utility leads to a higher within-group market share, causing reverse causality. A typical way of instrumenting this variable is using the number of products or rival suppliers in the group (e.g., Berry et al., 1995). This instrument is related to the within-group share in that more products or suppliers generally lead to lower within-group shares, whereas it does not affect the utility of the product itself. The second endogenous variable is the ticket price which is contained in the vector of product characteristics, $p_{ija} \in X_{ija}$. To instrument this variable we focus on cost shifters in the spirit of Hausman et al. (1994). These instruments are based on the idea that within-airline prices are correlated between markets due to common (shocks in) marginal costs, but after controlling for airline specific intercepts uncorrelated with market-specific valuation (see also, Nevo, 2000). We therefore can use prices charged in other similar markets, where we define similarity based on market distance brackets.

[3] We have tried different nesting structures, such as separate nests for LCC's and legacy carriers, but this did not lead to consistent results over all regions and consumer segments.

3.2. Data

This paper uses data from two main sources. OAG data provides information on passenger traffic, airfares and schedules. Missing fares in the US domestic market are augmented using the DB1B database published by the US DOT. The data set is completed with information from additional secondary sources, such as metropolitan population and gross domestic product (GDP) data from the Socioeconomic Data and Applications Center (SEDAC). To compute the potential demand d_{ij}, we use the geometric mean of the population residing in the two cities at the endpoints of a route.[4] We consider a catchment area of 50 km around each city. Accordingly, we compute the average GDP for each origin-destination pair using the geometric mean.

3.3. Sample selection

Markets are defined as directional one-way air travel between an origin and destination city in a given quarter. Thus, for air travel between London and New York in the first quarter of 2019, there is one market with London as the origin and New York as the destination, while there is another market the other way around. Origins and destination cities are defined on the city level so that a trip from London Heathrow (LHR) to John F. Kennedy International Airport (JFK) and a trip from London Gatwick (LGW) to Newark Liberty International Airport (EWR) both belong to the same market. This approach automatically accounts for overlapping airport catchment areas within multi-airport regions.

Products are defined as unique combinations of carrier groups and itineraries, where itineraries are distinguished by origin, destination and (if indirect) hub city.[5] We define carrier groups using the three major airline alliances and, for airlines outside these alliances, common ownership structures. Hence, the itinerary from LHR to JFK operated by British Airways is regarded as similar to the itinerary operated by American Airlines, as these carriers belong to the same alliance. Arguably, some of these aggregation steps lead to a loss of information, however, they also help alleviate the missing fare data problem by allowing fares to be averaged within products (see Appendix A).

The main estimates are obtained using the largest airline groups that together capture at least 95% of demand. This selection includes twelve airline groups in the European market, seven in the North American market and ten in the Transatlantic market. Excluding the smallest airlines makes it easier to fit and compute the demand models at the cost of little additional variation.

3.4. Parameter estimates

Table 1 presents the parameter estimates for the economy and business segments in the European, North American and Transatlantic markets respectively. The estimated coefficients and corresponding elasticities, measure the responsiveness of demand to changes in the key parameters of interest. In the base-case specification, air travel demand is affected by two market characteristics: (1) the geometric mean of the endpoints GDP per capita and (2) the market distance and market distance squared. Five product characteristics are also defined: the airfare, whether a flight is non-stop, frequencies, detour distance (as a proxy for travel time) and airline-specific constants (not shown in the table).

Market characteristics. We find the usual U-shaped curvature in air travel demand with respect to distance in the intra-continental European and North American markets, i.e. demand first increases in distance due to declining competition from other modes and then starts to decrease as distance increases further and air travel becomes more unpleasant. Consistent with the longer distances in the transatlantic market, distance has an immediate negative effect here. Somewhat surprisingly, the impact of distance in the transatlantic market turns positive at very long distances, but this is mostly an artefact of the non-random geographic placement of cities. The value of the coefficient is driven by some cities on the west coast of the US, such as Los Angeles and San Francisco, that generally have long market distances and a substantial amount of intercontinental air travel demand).

The impact of GDP per capita is positive in all six segments, indicating more travel between cities that are more affluent. In Europe and the transatlantic markets, the impact of GDP per capita is much larger for business travel than for economy, which one would expect given the greater number of business connections between wealthier cities and the idea that holiday destinations are not necessarily located in the wealthiest areas. Perhaps the latter is more the case in Europe than in North America, which could explain why a similar difference in the GDP effect on economy and business demand does not surface in the North American market.

Product characteristics. The parameters of the product characteristics are accurately estimated (low standard errors) and have the expected signs across all the markets. Consumer utility decreases with higher fares and increases with frequencies. Utility is also higher for non-stop flight products (relative to connecting products) and decreases with detour distance.

[4] This implicitly assumes symmetric demand between endpoint cities. Note that as markets are defined as directional one-way air travel (see next section) and given that the majority of air travellers book roundtrip travel, this assumption is generally well supported. In our data the deviation between demand from city A to B versus demand from city B to A is typically small, i.e. only a few percentage points relative to the directional one way demand, providing empirical support for the symmetric assumption.

[5] Consistent with much of the literature (Adler et al., 2010), we drop itineraries with more than two segments, so that there is a maximum of one hub city.

Table 1
Nested logit estimation results.

	Europe		North America		Transatlantic	
	leisure	business	leisure	business	leisure	business
Constant	– 9.80**	– 14.18**	– 9.38**	– 13.38**	– 10.21**	– 12.31**
	(0.06)	(0.11)	(0.06)	(0.05)	(0.13)	(0.15)
Market distance	0.54**	0.76**	0.65**	0.38**	– 0.13**	– 0.23**
(1,000 Km.)	(0.02)	(0.08)	(0.01)	(0.02)	(0.03)	(0.04)
Market distance^2	– 0.15**	– 0.10**	– 0.10**	– 0.04**	0.02**	0.02**
(1,000 Km.)	(0.01)	(0.02)	(0.002)	(0.004)	(0.002)	(0.002)
GDP per capita	0.40**	1.91**	0.61**	0.58**	0.68**	1.34**
	(0.04)	(0.11)	(0.03)	(0.05)	(0.04)	(0.05)
Directness	3.11**	3.10**	2.24**	2.55**	2.84**	2.65**
(0/1)	(0.03)	(0.03)	(0.02)	(0.02)	(0.03)	(0.02)
Detour distance	– 0.29**	– 0.13**	– 0.28**	– 0.21**	– 0.18**	– 0.13**
(1,000 Km.)	(0.01)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)
log frequency	0.45**	0.52**	0.42**	0.37**	0.34**	0.29**
	(0.004)	(0.02)	(0.003)	(0.01)	(0.005)	(0.01)
Ticket fare	– 0.60**	– 0.12**	– 0.46**	– 0.03**	– 0.18**	– 0.01
($100)	(0.03)	(0.02)	(0.03)	(0.01)	(0.01)	(0.004)
ρ	0.36**	0.22**	0.51**	0.31**	0.36**	0.24**
	(0.005)	(0.01)	(0.004)	(0.004)	(0.01)	(0.004)
Observations	89,081	27,869	183,727	42,905	82,905	36,682
Adjusted R^2	0.82	0.77	0.72	0.76	0.55	0.64

*Note: This table presents nested logit estimates, weighted by market size. Airline and quarter fixed effects are included in all models and coded using weighted-fixed effects coding (Sweeney and Ulveling, 1972), such that the constant in the utility function represents a weighted average over airlines and quarters. Standard errors are provided in parentheses. * Significant at the 5% level. ** Significant at the 1% level.*

In terms of magnitudes, economy demand is clearly more price sensitive.[6] The (own) price elasticity, which is the percentage change in demand if the price increases by 1%, averaged over all flight products, is equal to −1.19 in the European economy segment and −1.84 in the North American economy segment versus −0.72 and −0.32 in the European and North American business segments respectively. These elasticities demonstrate that the economy segment in North America is more price sensitive than in Europe, while for the business segment it is the other way around. Transatlantic passengers are less price sensitive in general, with average (own) price elasticities of respectively −0.80 and −0.23 for economy and business passengers, reflecting the lack of outside travel alternatives in this market.

What is also clear is that European travellers have a stronger dislike for indirect travel. If an indirect connection becomes direct, economy demand for this product increases on average by 488% in Europe and 456% in North America. For the business segment, this semi-elasticity of directness is 398% and 370% for Europe and North America, respectively.[7] The lower semi-elasticities for the business segments indicate that business travellers are somewhat less reluctant to fly indirectly, perhaps due to their travel experience and because indirect departure and/or arrival times may better fit their schedule (i.e., result in less schedule delay). Contrary to common understanding, we find higher elasticities for the economy segment vis-à-vis the business segment in terms of frequencies. The average own frequency elasticities in Europe and North America, respectively, are 0.69 and 0.83 for the economy segment versus 0.53 and 0.63 for the business segment. This counter-intuitive result might be due to remaining endogeneity in our model. In deregulated markets such as those studied here, airlines increase frequency in response to high demand, which makes the frequency endogenous to unobserved demand shocks. In the large and volatile economy segment, this may inflate the sensitivity of demand to frequency estimated in our model.[8]

In general, our empirical model fits air travel demand patterns in these regions quite well, which is a remarkable result given that we apply a single model specification to six regions and consumer segment combinations. This simplicity helps in applying these estimates in our game and also makes the estimates more easily adoptable by other researchers.

[6] Note that the (relative) magnitude of the effects (over columns) can not be easily grasped from the parameter estimates as these should be corrected for the inter-group correlation (ρ) and the scale of utility. Therefore it is better to consider the scale-free (own) elasticities, calculated as $[(1-m_{ijb(a)})(m_{a|ijb(a)}+\frac{1}{(1-\rho)}(1-m_{a|ijb(a)}))]\beta^x x_{ija}, \forall ija$ and where $x \in X_{ija}$ (see, Forinash and Koppelman, 1993), which we report in the text.

[7] While these semi-elasticities might be considered high, note that indirect travel alternatives generally have low passenger numbers and hence a high percentage increase equals a plausible increase in terms of absolute passenger numbers.

[8] One common way of addressing frequency endogeneity is to use the hub status of end point cities as an instrument (e.g. Berry and Jia, 2010). However, adding another instrument to our model, in addition to the instruments for the within-group market share and fare variables, leads to unstable estimates over the various subsegments considered here.

4. Game-theoretic formulation

In this section, we develop an applied game-theoretic formulation that models airline competition limited by slot constraints. Leveraging on the nested logit market share function discussed in Section 3, this setting replicates competitive interactions between airlines through a closed-form function. Decision variables for the airline players in the game include airfares, service frequencies, their fleet size and the number of seats per aircraft, on each link composing their relevant network.

We assume that all decision variables are continuous, which is common practice in strategic analyses. By treating variables as continuous, we leverage optimization processes and solution algorithms that are computationally efficient, enabling us to solve the model within a reasonable time-frame. While this assumption leads to small deviations from real-world outcomes given that some inputs are inherently discrete, the trade-off in terms of accuracy is relatively negligible compared to the significant gains in computational tractability.

We define networks, $G(\mathcal{N},\mathcal{K})$, in which each airport represents a node of the network belonging to the set $\mathcal{N}$. Legacy carriers utilize intercontinental hub-spoke networks whereby the hubs are connected to the spokes through ordered legs within the set $\mathcal{K}$. For simplicity, we assume that the airlines do not code-share or interline and we discuss the implications of relaxing this assumption in the case study. LCCs serve one continent using a single aircraft type and a fully connected network. Whilst legacy carriers offer flights on itineraries connecting an origin to a destination city either directly or indirectly, LCCs only offer service on each individual leg. A direct flight, denoted by $\delta = 1$, connects two points in the network without a stop. Assuming all other trip characteristics are equal, passengers tend to prefer direct and shorter flights when choosing an airline and itinerary. Therefore, we quantify the additional time and distance involved in connecting itineraries compared to direct flights for the same origin-destination pair. Moreover, connecting flights involve additional time due to the need for a feasible connection at the hub airport between the two segments of the indirect trip.

We now define four subsets of itineraries and legs along with two subsets related to the slot constrained hubs, in order to streamline the identification of the constrained nodes and arcs within the network.

$$\mathcal{N}^\circ = \{i^\circ, j^\circ | i^\circ, j^\circ \text{ are the itineraries passing through leg } k \in \mathcal{K}\}$$
$$\Omega = \{\omega | \omega \text{ is the leg directly connecting } i, j \in \mathcal{N}\}$$
$$\Omega' = \{\omega' | \omega' \text{ is the first leg of the itinerary } i, j \in \mathcal{N}\}$$
$$\Omega'' = \{\omega'' | \omega'' \text{ is the second leg of the itinerary } i, j \in \mathcal{N}\}$$
$$\hat{\mathcal{N}} = \{\hat{i} | \hat{i} \text{ are the slot constrained nodes}\}$$
$$\Upsilon = \{\upsilon | \upsilon \text{ are the legs connecting a slot constrained node } \hat{i} \in \hat{\mathcal{N}}\}$$

Airlines operate at airports by conducting landings and take-offs within designated time banks $\bar{\upsilon}$. When slot availability is reduced, the carriers must adjust by decreasing flights and redistributing operations across these banks. Assuming slot reductions are evenly distributed throughout the day, maintaining an existing connection results in additional layover time $\hat{\upsilon}$, effectively extending the bank duration. Fig. 1 illustrates how slot reductions impact layover times within the bank structure. We formally define the additional delay at slot constrained airports as follows:

$$\hat{\upsilon}_{ija} = \begin{cases} \bar{\upsilon}_{ija}\left(1 - \varrho_i\right) & \text{if} \quad i \in \hat{\mathcal{N}} \vee j \in \hat{\mathcal{N}} \\ 0 & \text{otherwise} \end{cases} \qquad \forall i, j \in \mathcal{N}, a \in \mathcal{A} \tag{5}$$

where $\varrho_{\hat{i}} \in (0, 1]$ represents the proportion of slots remaining at the constrained hub $\hat{i}$, leading to a proportional increase in layover time corresponding to the percentage reduction in slots for airline $a \in \mathcal{A}$.

A reduction in available slots reduces passenger utility by increasing layovers, represented as an additional detour distance $\tilde{g}_{ija}$. The detour distance is defined as the difference between the sum of the distances of connecting flights in an itinerary and the direct flight distance between the origin and destination. By using the average aircraft speed $\bar{s}_h$ per type of aircraft h, the extra layover time is translated into an equivalent distance, representing the added travel time imposed by slot constraints at the affected airport. The final detour distance for flights to and from a slot-constrained airport, including the extra layover time, is given by:

$$\tilde{g}_{ija} = \begin{cases} (g_{\omega'} + g_{\omega''}) - g_{\omega} + \hat{\upsilon}_{ija}\bar{s}_h & \text{if} \quad (\delta = 0) \wedge (i \in \hat{\mathcal{N}} \vee j \in \hat{\mathcal{N}}) \\ \bar{g}_{ija} & \text{otherwise} \end{cases} \qquad \forall i, j \in \mathcal{N}, h \in \mathcal{H}, a \in \mathcal{A} \tag{6}$$

European LCCs often operate from airports in peripheral locations to reduce operational expenses. To account for this, we incorporate an additional term in the utility function to represent airport access time. We assume that passengers on European LCC flights require 3 h to access and egress the remote airports, whereas only 1.5 h is required to access European legacy carrier services and all carrier services in North America. Analogous to the approach above, this extra travel time is converted into an equivalent flight detour for passengers using low-cost services.

In the remainder of this section, we present the components of the model including operating costs, ownership costs, the profit function and constraints. Then we outline the game and algorithm applied, concluding with the social welfare function used to compare multiple scenarios. Table 2 summarizes the notation used.

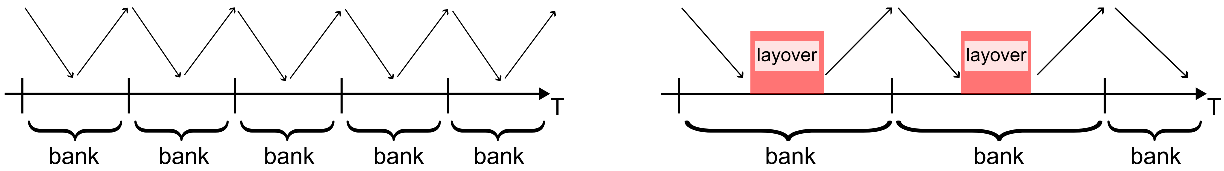

(a) Business as usual (b) Slot constrained

Fig. 1. Bank structure with and without slots constraints.

Table 2
Notation.

Sets and Indices		
$\mathcal{A}$	Set of airlines; indexed by a	
$\mathcal{H}$	Set of flights (i.e. short and long haul); indexed by h	
$\mathcal{K}$	Set of all legs in the network; indexed by k	
$\mathcal{N}$	Set of airport nodes; indexed by i, j	
$\mathcal{T}$	Set of passenger types; indexed by t	
Λ	Set of local pollutants; indexed by λ	
Parameters		**Unit**
d_{ij}	Potential demand for itinerary (i, j)	$passengers$
β_{xt}	Utility parameter per market attribute x and passenger type t	$utility/attribute$
ρ_t	Heterogeneity between nests in discrete choice model per passenger type t	-
c_{ka}	Cost to serve leg k for airline a	$€/flight$
δ_{ija}	$=1$ if connection between i and j operated by a is direct; 0 otherwise	-
$\bar{f}_h$	Average quarterly utilization of aircraft by flight type h	$flights/quarter$
$\underline{f}_{\hat{i}a}$	Movements at airport i in the unconstrained case for airline a	$flights/quarter$
g_k	Great circle distance of leg k	km
$\tilde{g}_{ija}$	Detour distance between nodes i and j for airline a inclusive of extra time $\tilde{\upsilon}$	km
ϕ_h	Fuel per kilometer by aircraft type h	t_{fuel}/km
$\bar{p}$	Fuel price in dollars per ton of kerosene	$\$/t_{fuel}$
o_h	Ownership cost for aircraft type h	$€/aircraft$
co_2	Conversion factor between fuel consumption and co_2 produced	t_{CO_2}/t_{fuel}
χ	Share of operating costs without fuel and ownership costs	$\%$
Ψ_h	Purchase price of aircraft of type h	€
σ_{ia}	Maximum slots for airline a at airport i	$flights$
n	Aircraft life in years	$years$
τ_h	Salvage value of aircraft type h in period n	€
ℓ	Interest rate	$\%$
ε_k	Tons of CO_2 generated per km on leg k	t_{CO_2}/km
ζ	Social cost of carbon	$€/t_{CO_2}$
η_k	Social cost of noise on leg k	$€/pkm$
$\bar{s}_h$	Average aircraft speed per type h	$km/hour$
ϱ_i	Percentage of slots remaining at airport i	$\%$
$\bar{\upsilon}_{ija}$, $\hat{\upsilon}_{ija}$	Bank hours pre and post slot constraints for a connection on itinerary ij	$hour$
$U_{f_{ka}}$	Upper bound on frequency per leg k per airline a	$flights$
U^s_k	Upper bound on seats offered on leg k	$seats$
Decision variables		
f_{ka}	Service frequency on leg k of airline a	$flights$
p_{ijta}	Fare for itinerary (i, j) per passenger type t on airline a	€
$\tilde{\alpha}_{ha}$	Number of aircraft owned of type h by airline a	$aircrafts$
s_{ka}	Seats available on leg k served by airline a	$seats$
Auxiliary variables		
$\hat{m}_{ijta}$	Market share for itinerary (i, j) and passenger type t of airline a	$\%$
V_{ijta}	Systematic utility for itinerary (i, j) per passenger type t of airline a	$utility$
z_{ija}	Minimum frequency over indirect itinerary (i,j) for airline a	$flights$

4.1. Cost function

According to Swan and Adler (2006), the direct operating costs of an airline are a function of great circle distance, g_k, the number of seats, s_k, on an aircraft and whether the aircraft is short-haul narrow-body or long-haul wide-body. The functional form provided by Swan and Adler (2006) for different aircraft categories typically involves terms such as $(g_k + A)(s_k + B) \times P_{\text{original}}$, where A and B are constants related to distance and seat capacity and P_{original} is a cost coefficient. For our analysis, we adopt this structure. The specific constant terms A and B are taken as $A = 722$, $B = 104$ for short-haul narrow-body aircraft ($k \in \mathcal{K}^s$) and $A = 2200$, $B = 211$ for long-haul wide-body aircraft ($k \in \mathcal{K}^l$), consistent with the parameters for these categories in the Swan and Adler (2006) framework.

The original cost coefficients (P_{original}) from Swan and Adler (2006) were presented in US dollars and at 2006 cost levels. To update these coefficients for our study to 2019 euro values, we performed the following modifications: (i) the original US dollar cost coefficients were multiplied by the average inflation growth of 1.2 to translate the 2006 costs to 2019 levels; (ii) the resulting 2019 US dollar values were then converted to euros. We used an average 2019 exchange rate of \$1 = €0.8931 for this conversion. These two adjustments applied to the original coefficients yield the updated cost parameters used in our model: €0.0255 for short-haul aircraft and €0.0154 for long-haul aircraft.

In the original cost function, operating costs and aircraft ownership costs are combined. In our optimization framework, however, ownership costs are modelled explicitly as a function of fleet size which is a decision variable. We therefore disaggregate the cost function by removing the ownership cost component. According to Swan and Adler (2006), ownership costs represent approximately 30% of total operating costs. Accordingly, we set $\chi = 0.70$ to exclude this share from the operating cost function, ensuring that ownership costs are accounted for only once in the objective function. The resulting direct operating cost function, excluding ownership, is:

$$c_k = \begin{cases} \chi(g_k + 722)(s_k + 104)€0.0255 & \text{if} \quad k \in \mathcal{K}^s \\ \chi(g_k + 2,200)(s_k + 211)€0.0154 & \text{if} \quad k \in \mathcal{K}^l \end{cases} \tag{7}$$

where

$$\mathcal{K}^s = \{k | k \in \mathcal{K} \text{ are the short-haul flights served}\}$$

$$\mathcal{K}^l = \{k | k \in \mathcal{K} \text{ are the long-haul flights served}\}$$

We account for the distinct operating models of LCCs by adjusting their operating costs. These are modelled as 80% of the costs c_k incurred by legacy airlines in North America and 60% in Europe, respectively. These adjustments are based on comparative analyses of Costs per Available Seat Kilometre (CASK) derived from the airlines' annual reports.

The quarterly cost of owning an aircraft is approximated by the equivalent annual capital costs divided accordingly:

$$o_h = \frac{1}{4}\left[(\Psi_h - \tau_h)\left(\frac{\mu\ell(1+\ell)^n}{(1+\mu\ell)^n - 1}\right) + \tau_h \mu\ell\right] \tag{8}$$

where Ψ_h is the initial purchase price of an aircraft of type h; τ_h is the salvage value at the end of the n-year time period; and ℓ is the interest rate approximated by the Weighted Average Cost of Capital (WACC) in 2019, set to be 10% (Iata, 2019). We have selected two of the most commonly purchased aircraft models as reference models for the aircraft types (see Table 3).[9] The purchase price of these aircraft are based on the average list prices.[10] Salvage values are derived by assuming straight-line depreciation over 30 years, with a service life as a passenger aircraft of 20 years. Therefore, the salvage value is equal to one-third of the purchase price. We further assume a discount factor on the purchase price of an aircraft of 25%. Discounts are common practice in the aviation industry (Pulvino, 1998).

4.2. Profit function

Airlines maximize profits, π_a, by choosing service frequency, f_{ka}, aircraft capacity, s_{ka} and airfares p_{ijta} in addition to the number and type of aircraft $\tilde{\alpha}_{ha}$ operating over their chosen network.[11] Given this setting, the objective function for airlines is modelled as follows:

$$\max_{\substack{p_{ijta}, f_{ka}, \tilde{\alpha}_{ha}, \\ s_{ka}, \hat{m}_{ijsa}}} \pi_a = \sum_{\substack{ij \in \mathcal{N} \\ i \neq j}} \sum_{t \in \mathcal{T}} d_{ij} \hat{m}_{ijta} p_{ijta} - \sum_{k \in \mathcal{K}} c_k f_{ka} - \sum_{h \in \mathcal{H}} o_h \tilde{\alpha}_{ha} \tag{9}$$

where $\hat{m}_{ijsa}$ is an auxiliary variable that denotes the market share function specified in Eq. (3) representing the share of demand for each city pair *(i,j)* per passenger type served by airline a, c_k, represents the operating costs, defined in Eq. (7), incurred by the airline to serve a specific leg, o_h is the quarterly ownership cost of aircraft type h (either long-haul wide-body or short-haul narrow-body) and $\tilde{\alpha}_{ha}$ is the number of aircraft of type h that the carrier deploys across its network. We note that service frequency influences both the revenue function (by increasing passengers' willingness to pay) and the cost function (by raising operating expenses and fleet requirements), thus creating a trade-off for airlines.

We now proceed to define the constraints of the optimization problem as follows:

$$\hat{m}_{ijta} = \frac{\exp\left(\frac{V_{ijta}}{1-\rho_t}\right)}{\sum_{a' \in \mathcal{A}} \exp\left(\frac{V_{ijta'}}{1-\rho_t}\right)} \frac{\left[\sum_{a' \in \mathcal{A}} \exp\left(\frac{V_{ijta'}}{1-\rho_t}\right)\right]^{(1-\rho_t)}}{1 + \left[\sum_{a' \in \mathcal{A}} \exp\left(\frac{V_{ijta'}}{1-\rho_t}\right)\right]^{(1-\rho_t)}} \quad \forall\, i,j \in \mathcal{N}, t \in \mathcal{T} \tag{10}$$

$$z_{ija} \le f_{\omega' a} \quad \forall\, i,j \in \mathcal{N}, \omega' \in \Omega' \tag{11}$$

$$z_{ija} \le f_{\omega'' a} \quad \forall\, i,j \in \mathcal{N}, \omega'' \in \Omega'' \tag{12}$$

9 Using competing manufacturers' models does not lead to materially different capital cost values.

10 Published on Axon Aviation (2022).

11 Although the network structure is fixed, airlines may discontinue a route by setting frequency to zero. This provides carriers with the flexibility to choose the markets to serve.

$$\sum_{i^\circ j^\circ \in \mathcal{N}_k^\circ} \sum_{t \in \mathcal{T}} d_{i^\circ j^\circ} \hat{m}_{i^\circ j^\circ t} \leq s_{ka} f_{ka} \quad \forall\ k \in \mathcal{K} \tag{13}$$

$$\sum_{k \in \mathcal{K}} f_{ka} \leq \tilde{\alpha}_{ha} \bar{f}_h \quad \forall\ h \in \mathcal{H} \tag{14}$$

$$s_{ka} \leq U_k^s \quad \forall\ k \in \mathcal{K} \tag{15}$$

$$f_{ka} \leq U_{f_{ka}} \quad \forall\ k \in \mathcal{K}, h \in \mathcal{H} \tag{16}$$

$$\sum_{v \in \Upsilon_{\hat{i}}} f_{va} \leq \sigma_{\hat{i}a} \forall\ \hat{i} \in \hat{\mathcal{N}} \tag{17}$$

$$f_{ka} \geq 0, s_{ka} \geq 0, \quad \forall\ k \in \mathcal{K} \tag{18}$$

$$p_{ijta} \geq 0, \quad \forall\ i, j \in \mathcal{N},\ t \in \mathcal{T} \tag{19}$$

$$\tilde{\alpha}_{ha} \geq 0, \quad \forall\ h \in \mathcal{H} \tag{20}$$

Constraint (10) specifies the market share function using the nested logit formulation from Eq. (3). Eqs. (11) and (12) introduce a linear disaggregation of the minimum function, representing the bottleneck frequency in connecting itineraries. In the passengers' utility function, the bottleneck frequency is defined as the minimum of the two connecting flight frequencies, creating a discontinuity. In order to address this, we reformulate the problem using linear disaggregation methods to decompose the minimum function. This reformulation is advantageous because the linear constraints do not require additional binary variables or big-M formulations, as is typical in general disjunctive programming. Eq. (13) ensures that seat capacity on a specific leg is equal to or greater than the total demand served across all itineraries using that leg. Constraint (14) ensures a sufficient fleet given average utilization for long-haul and short-haul routes. Eq. (15) sets upper bounds on the number of seats offered per leg according to the type of aircraft employed and Eq. (16) ensures that airlines meet slot restrictions. Eq. (17) enforces the slot constraint per airline per constrained hub by defining $\sigma_{\hat{i}a} = \underline{f}_{\hat{i}a} \varrho_{\hat{i}}$ as the movements allowed when slot constrained, where $\underline{f}_{\hat{i}a}$ is the sum of movements for each airline at airport $\hat{i}$ in the unrestricted case. This constraint is applied only when required by the specific scenario. Constraints (18) to (20) ensure that all decision variables are non-negative.

4.3. Game-theoretic competition and algorithm

The single-stage game involves a set of airlines, $\mathcal{A}$, whose best response function (Eq. (9)) enable them to make decisions while considering the choices of their competitors in a complete and perfect information setting. The algorithm outlines the solution method for this simultaneous-move game. It begins with initial values and the network setup, then iteratively solves the optimization problem for each airline until a cycle is completed. The process continues until two consecutive cycles yield results within a predefined threshold of 1%, signalling convergence to a Nash-equilibrium. The optimization problem consists of a non-linear objective function with non-linear constraints, which is solved using *IPOPT* (Wächter and Biegler, 2006), an interior-point solver for non-convex, continuous problems (the algorithm is described in Appendix B). Since the global optimality of the obtained solution cannot be guaranteed, we assess the robustness of our results by varying the initial starting points and the ordering of players. We compute the relative deviation in solutions across fifty runs initialized with random starting points and randomized player orderings. The results of the model are robust to variations in initialization, with relative changes in the solution not exceeding 0.02%, as illustrated in Fig. B.1.

4.4. Social welfare

To evaluate the wider economic impacts of slot reductions at a hub, we assess the changes in social welfare before and after the reductions are implemented. Given the potentially significant effects from slot reductions on airline profitability and passenger connectivity, particularly at major airports, we define social welfare as a function of consumer surplus, producer profits and environmental externalities:

$$w = \sum_{ij \in \mathcal{N}} \sum_{t \in \mathcal{T}} \frac{d_{ij}}{-\beta_{price}} \ln \left| 1 + \left(\sum_{a' \in \mathcal{A}} \exp\left(\frac{V_{ijta'}}{1-\rho_t} \right) \right)^{1-\rho_t} \right| + \sum_{a' \in \mathcal{A}} \pi_{a'} \left(f^*_{ka'}, p^*_{ijta'}, \tilde{\alpha}^*_{ha'}, s^*_{ka'} \right) - \sum_{a' \in \mathcal{A}} \sum_{k \in \mathcal{K}} \varepsilon_k \zeta f^*_{ka'}$$
$$- \sum_{a' \in \mathcal{A}} \sum_{k \in \mathcal{K}} \sum_{\lambda \in \Lambda} \epsilon_{k\lambda} \xi_\lambda f^*_{ka'} - \sum_{a' \in \mathcal{A}} \sum_{k \in \mathcal{K}} \eta_k f^*_{ka'} \tag{21}$$

where $\varepsilon_k = g_k \phi_h co_2$ is the CO_2 produced on flight leg k given aircraft fuel efficiency ϕ_h, ζ is the social cost of a ton of carbon, $\epsilon_{k\lambda}$ is the level of local pollutant, λ, generated per landing and take-off (LTO) cycle multiplied by the respective social cost ξ_λ and η_k is the social cost of noise for a movement on leg k. In order to account for other non-CO_2 global pollutants, we convert the CO_2 produced per flight into CO_2-equivalents by multiplying by three, as suggested in Lee et al. (2021). In Eq. (21), the first row computes the consumer surplus based on the log-sum of passenger utility (Small and Rosen, 1981); the second row sums the profits of all airlines less the sum of network-wide global emissions; and the third row computes the costs associated with local pollutants emitted and noise caused by aircraft movements.

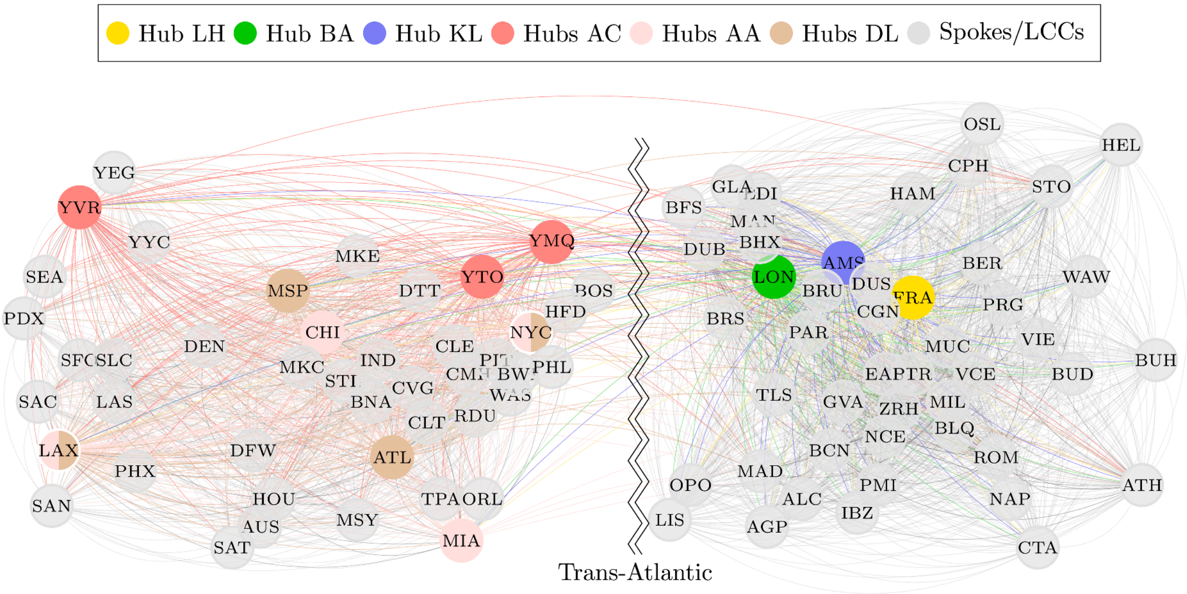

Fig. 2. Network covering North American and Western European continents.

5. Case study: slot restrictions in North America and Europe

In this section, we apply our game-theoretic formulation on a large-scale network covering North America (NA) and Europe (EU) as shown in Fig. 2. This network includes the intra-continental markets in both regions as well as the intercontinental Transatlantic (TRA) market. The network consists of 91 nodes, 47 in Europe and 44 in North America, selected from among the 50 busiest airports in each region.[12] Coloured nodes correspond to hub airports for the legacy carriers, while gray nodes represent spoke airports that are served either by legacy carriers or by low-cost carriers operating within that region.

The nodes in our network collectively account for 44% of the quarterly demand across the three markets. The average distances between nodes reflect real world conditions, namely an average flight of 1,225 km in the EU, 1,935 km in NA and 7,456 km for TRA routes. TRA connections are served by legacy carriers operating through their hubs, with the EU market structured as pure hub-and-spoke networks, whereas the NA market features multi-hub-and-spoke configurations.[13] Conversely, LCCs operate in a fully connected network within the region where they are based.

Following the OAG classification, we treat the following airlines as LCCs: Ryanair (FR), easyJet (U2), Eurowings (EW), Wizz Air (W6) and Vueling (VY) in Europe and Southwest (WN), JetBlue (B6) and Spirit (NK) in North America; while Lufthansa (LH), British Airways (BA), KLM (KL), Air Canada (AC), American Airlines (AA) and Delta Air Lines (DL) are treated as legacy carriers. Our model assumes that these airlines compete independently, with no agreements such as alliances or joint ventures in place. We relax this assumption in a sensitivity analysis, reflecting cross-continental joint ventures operating across the entire network and show that the key dynamics of our model remain intact (D).

We assume that airlines operate a fleet based on the most commonly used aircraft: the Boeing 737 for short-haul flights and the Boeing 777 for long-haul flights. As a robustness check, in Appendix E we incorporate the impact of fleet renewal induced by a carbon tax (De and G, 2022) on our results. For short-haul flights, an upper bound on seat capacity is set at $U_k^s = 168$ for legacy carriers and $U_k^s = 187$ for LCCs. For long-haul flights, the seat capacity has an upper bound of $U_k^s = 370$. The upper bound on flight frequency, $U_{f_k a}$, for each leg and airline is derived from the OAG schedule data. Average speeds, including the LTO cycle, are 561 km/h for short-haul flights and 775 km/h for long-haul operations. Regional differences in aircraft utilization are accounted for, with an average of 3 flights per day in the EU and 2 flights per day in NA, reflecting variations in stage lengths. For simplicity, we assume a maximum load factor of 80% for all carriers. Each time bank is assumed to be 4 h for landings and take-offs, which ensures sufficient layover time for connecting intercontinental passengers. Aircraft purchase and salvage values, used to compute ownership costs, are listed in Table 3.

The social cost of a ton of CO_2 is set at € 200, based on the latest IPCC estimates (Pörtner et al., 2022). Local emissions generated during landing and take-off cycles are calculated using the Master Emission Calculator (Eea, 2023), with their health impacts assessed by the European Environment Agency (Eea, 2012). These values are summarized in Table 4. Additionally, we compute two average

[12] We removed some airports that created imbalances in expected average stage lengths (e.g., remote airports in Alaska and Hawaii).

[13] Passengers select routes through hub(s) offering the shortest travel time between origin and destination cities among all possible itineraries per airline.

Table 3
Aircraft purchase and salvage values.

h	Reference model	Ψ (€ M)	τ (€ M)
Short-haul	B737	98.24	29.77
Long-haul	B777	330.45	100.14

Table 4
Local pollutants and their social cost.

Pollutant (λ)	Social cost (€/tonne)	Quantity (kg/LTO)	
		B737	B777
HC	7288	1.64	3.93
PM	237,123	0.07	0.20
NOx	42,953	11.28	68.23
SOx	38,345	0.69	2.44

Table 5
Base-case airline results in quarterly values.

	European		American	
	Legacies	LCCs	Legacies	LCCs
Pax. short econ.	7,742,917 (8,961,543)	28,827,918 (38,099,656)	28,555,827 (34,222,650)	26,464,235 (31,073,389)
Pax. short bus.	752,360 (919,728)	–	739,878 (1,939,391)	–
Pax. long econ.	1,129,608 (1,556,390)	–	1,534,612 (1,965,689)	–
Pax. long bus.	83,358 (317,940)	–	132,084 (160,194)	–
CASK (¢)	7 (7)	3 (4)	7 (7)	4 (5)
RASK (¢)	15 (8)	8 (5)	14 (8)	8 (6)
Fare short econ. (€)	183 (140)	149 (113)	242 (195)	217 (178)
Fare short bus. (€)	846 (464)	–	3070 (786)	–
Fare long econ. (€)	509 (450)	–	583 (425)	–
Fare long bus. (€)	8434 (3,575)	–	8819 (3,363)	–
Freq. short	85,966 (155,499)	236,765 (259,610)	265,101 (836,579)	179,158 (317,688)
Freq. long	3855 (16,611)	–	5273 (25,086)	–

Note: In brackets we report the real-world values. “Short” refers to domestic and regional flights within the continent; “long” to transatlantic flights. “Econ.” indicates economy class, while “bus.” denotes business class passengers.

airport noise costs per passenger-kilometre (pkm) for short-haul and long-haul operations of 0.46 and 0.01 $€\text{-}cent/pkm$, respectively, following the European Commission (2020) handbook.

In line with the estimation process described in Section 3, we consider all flights as one-way oriented trips. This implies that all frequencies and slot valuations refer to a single movement, which may consist of a landing or a take-off. In practice, airline slots are allocated over an extended time horizon, typically a season. Our aim is to compute the marginal valuation arising from a relative change in slot availability, hence we focus on a single slot defined as either a landing or a take-off. This approach enables a more granular assessment by capturing differences between long- and short-haul slot usage.

We first solve the base-case scenario for all airlines without additional slot restrictions and the Nash-equilibrium outcomes are reported in Table 5. For the most part, results align well with real-world outcomes from 2019, shown in brackets. However, Revenue per Available Seat Kilometre (RASK) values exceed those reported by the airlines because the nodes in our case study focus primarily on major metropolitan areas, which typically represent denser and more profitable destinations. Furthermore, predicted long-haul frequencies are lower than observed real-world values, due to the absence of connecting passengers originating from nodes external to the network modelled.

5.1. Slot reduction scenarios

We now solve the model under multiple slot reduction scenarios for all major cities in the network. These reductions apply to all carriers operating at the constrained hubs. Slot availability is reduced incrementally from 0 to 20% of the airline’s 2019 service frequencies at the hub, in steps of 5%. The results presented below are specific to the Europe-North America network used in our case study and should be interpreted relative to that network. This qualification does not affect the generality of our methodology, which can be applied without modification to alternative network definitions.

Overall social welfare declines as slot constraints increase, as shown in Table 6 and Fig. 3. The profitability of most carriers is minimally affected by slot reductions because of three compensating effects. Slot restrictions reduce passenger volumes, but fewer flights also lower operating costs. Moreover, equiproportional reductions across all carriers weaken competition and increase market power, permitting higher fares. As described in Appendix C, the price increase is relatively higher in Europe than in North America, as the multi-hub structure in North America reduces passengers’ exposure to slot constraints. Airlines also tend to eliminate marginal,

Table 6
Average social welfare in € M and percent variations for its components in EU and NA.

	Slot reduction (%)				
	0	−5	−10	−15	−20
Europe					
Consumer surplus	8572.35	−0.47%	−0.95%	−1.46%	−1.99%
Producer surplus	3781.01	0.18%	0.32%	0.42%	0.49%
Global pollution	−2366.78	−0.55%	−1.12%	−1.70%	−2.28%
Local pollution	−362.85	−0.74%	−1.48%	−2.22%	−2.96%
Noise	−132.83	−0.51%	−1.03%	−1.57%	−2.12%
Welfare	9490.90	−0.18%	−0.38%	−0.62%	−0.89%
North America					
Consumer surplus	16,322.79	−0.40%	−0.82%	−1.24%	−1.67%
Producer surplus	6945.03	0.26%	0.49%	0.68%	0.82%
Global pollution	−4806.50	−0.37%	−0.77%	−1.19%	−1.62%
Local pollution	−513.52	−0.51%	−1.02%	−1.54%	−2.05%
Noise	−253.98	−0.35%	−0.74%	−1.14%	−1.56%
Welfare	17,693.82	−0.15%	−0.31%	−0.49%	−0.69%

low-yield routes first, often short-haul services to secondary markets, so the loss in traffic is more easily offset by the ability to raise fares elsewhere. These effects largely cancel out, resulting in small profit changes even with a 20% reduction in slots, with producer surplus increasing slightly.[14]

Consumer surplus, which constitutes the largest share of total welfare, declines substantially as connectivity reduces while fares increase due to greater airline market power. This decline reflects the loss of connectivity, convenience and affordability experienced by passengers. It should be noted that the relative changes in consumer surplus may appear small because, coherently with our estimation process and the fact that externalities have broader impacts, we are considering the consumer surplus of all potential demand in the surrounding region of our network cities, which includes individuals that are not travelling and whose surplus is not affected (i.e., there is a large constant).

In contrast to the change in consumer surplus, the environmental impacts of slot reductions contribute positively to social welfare. Slot constraints reduce flight operations, leading to lower fuel consumption and emissions, as well as reduced noise at airports. However, we find that the welfare gains from reduced global and local emissions as well as noise pollution, the key motivations for reducing slot allocations, are relatively modest compared to the impact on consumer surplus.

The cumulative values of the individual welfare components are illustrated in Fig. 4. Consumer surplus remains the largest welfare component but experiences a sharp decline as slot reductions intensify, reflecting the loss of connectivity for passengers. Conversely, airlines may initially increase profits because an equiproportional reduction in slots across carriers weakens competition at constrained airports. Under severe restrictions, however, revenue losses from lower passenger volumes erode these market power gains (see Appendix F). Finally, slot-constraint policies yield a steadily increasing reduction in negative externalities, most notably through global emissions savings.

We now examine marginal slot values to understand how incremental changes in slot availability influence societal outcomes. As shown in Fig. 5, slot valuations increase across all cities as constraints tighten, reflecting rising marginal social costs. The average marginal social value per slot ranges from €10,900 to €13,140 as reductions increase from 5 to 20%. The results also reveal heterogeneous effects across cities. Three distinct clusters emerge: Canadian airports show the highest marginal values, US airports fall in an intermediate range and European airports register the lowest. This pattern suggests that geographic location and network position are key drivers of slot value. Canadian cities, located at the edge of the network with limited connectivity and lower traffic volumes, place greater value on their limited slots.[15] Conversely, centrally located hubs with higher volumes of connecting traffic, particularly London, show consistently lower marginal values. European airports rank lower overall because smaller passenger volumes combine with environmental savings from reduced long-haul operations to diminish slot values. Yet even in London, where marginal social value is lowest, the environmental benefits and modest gains in producer surplus do not outweigh the welfare losses from reduced connectivity and higher fares.

In Fig. 6 we present the individual components of the marginal slot values. From the consumers' perspective, the value increases as constraints tighten, reflecting the growing welfare loss from reduced connectivity and higher fares. For airlines, the marginal value of a slot is consistently negative, indicating that they gain from the fact that competition is limited by uniformly reducing slot numbers. This strategy boosts their market power and enables them to charge passengers higher fares. This effect is especially marked in North America, where more negative valuations indicate greater surplus extraction from passengers. However, as slot availability further declines, the incremental profit benefits diminish, revealing a trade-off between increased pricing power and reduced passenger

[14] In Appendix F, we solve the game for slot reductions exceeding 20% (50 and 70%), showing that airline profits eventually decline under severe restrictions. However, such capacity reductions are unrealistic in practice.

[15] Hubs such as Los Angeles and Vancouver are central to North America-Asia markets, which are not included in the current network.

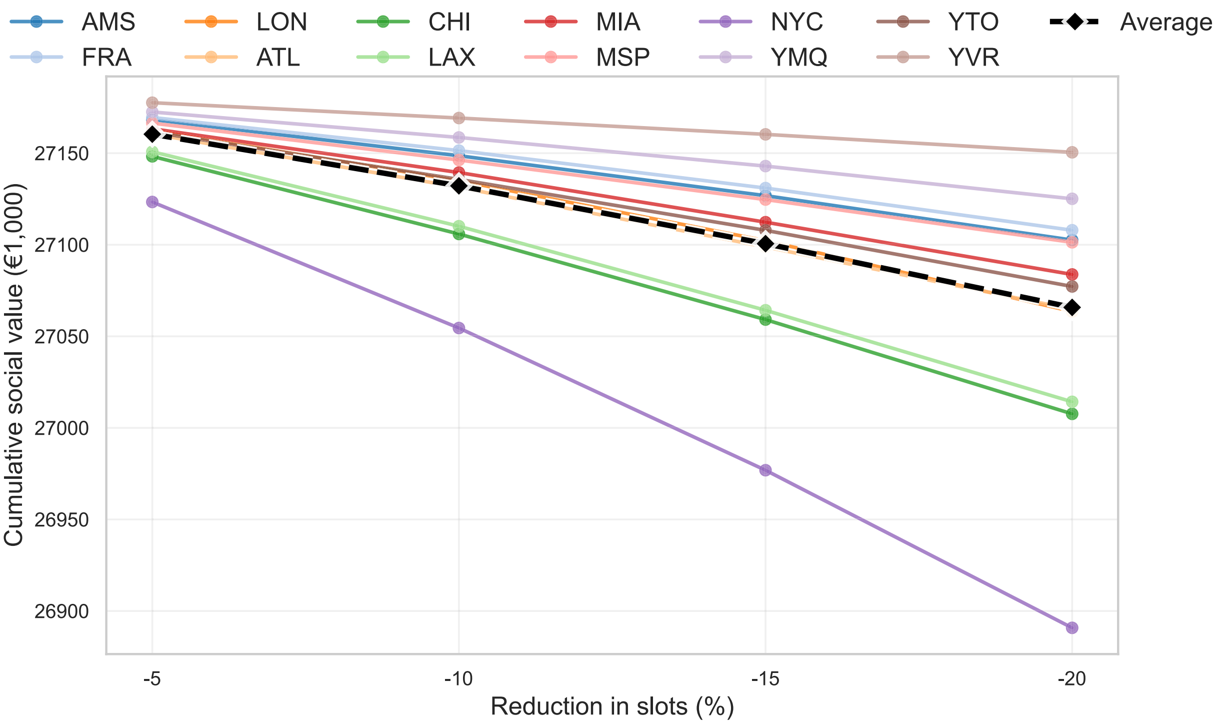

Fig. 3. Cumulative social value of a slot.

Note: This figure reports the cumulative impact of slot reductions on overall social welfare, both at individual hubs and on average across the Europe-North America network.

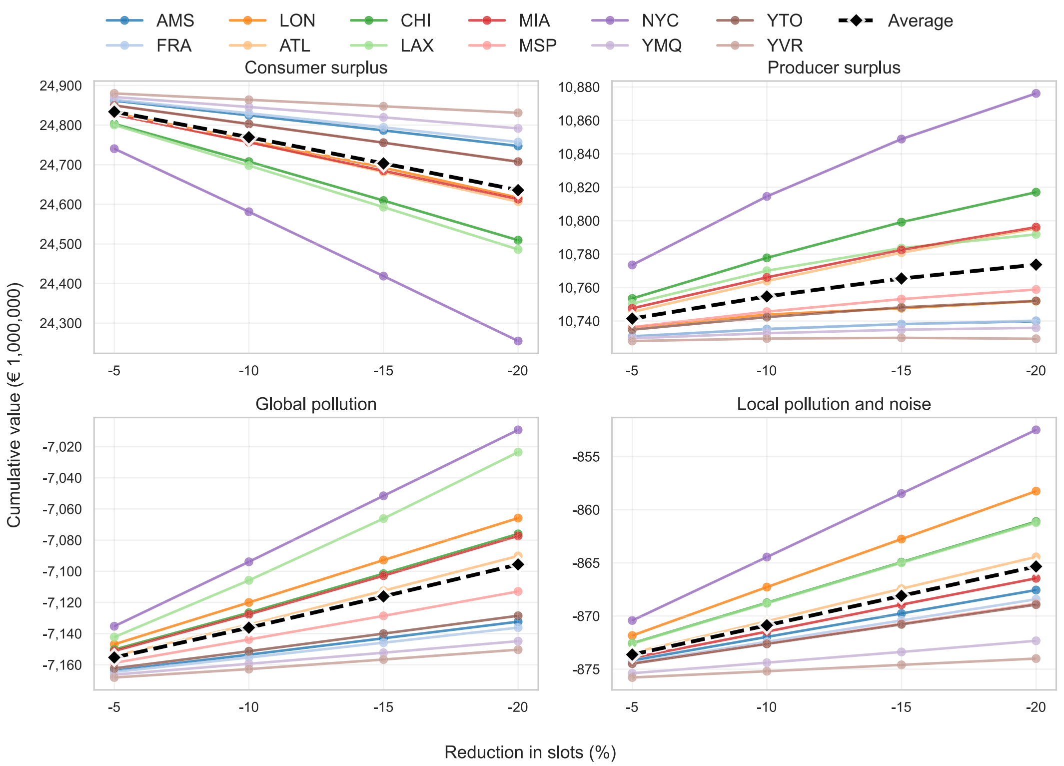

Fig. 4. Cumulative value of welfare components as a function of slot reductions.

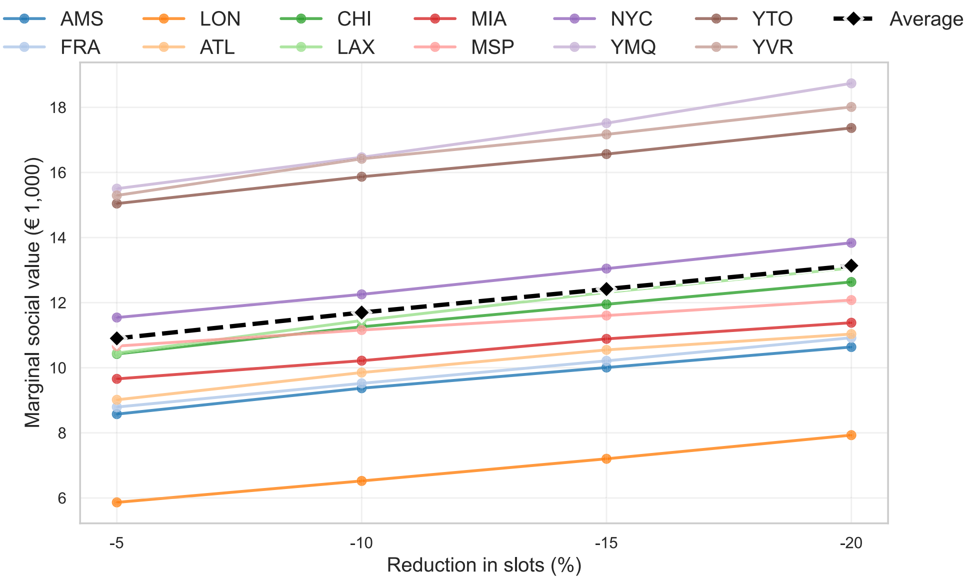

Fig. 5. Marginal social value of a slot.

volumes, evident in the upward trend of airline valuations in Fig. 6.[16] From an environmental perspective, as supply decreases, both global and local externalities decrease accordingly. Whilst local pollution and noise have relatively minor impacts, the marginal reductions in greenhouse gas emissions, contrails and other non-CO_2 emissions play a more important role.

At the airline level, as shown in Fig. 7, the implementation of slot constraints leads to a notable reduction in short-haul flights, whereas long-haul flights experience a relatively minor cut compared to the initial allocation of short and long-haul operations. This outcome is driven by the strategic prioritization of long-haul operations by airlines, which typically yield higher profits compared to short-haul destinations.

Fig. 8 shows the marginal slot valuation by carrier type and geographic relation to the constrained airports. In the top-left panel, airlines based at constrained hub airports show the most negative valuations under mild slot reductions. As constraints tighten, these valuations become less negative and even turn positive in some cases, such as British Airways, as declining passenger volumes through the hub erode the profitability of long-haul routes and thereby reduce the gains from market dominance. This demonstrates that the legacy carrier at the hub where slot are reduced initially has the most to gain from reduced competition, but that as slot reductions become more severe, they also are the first to lose. The top-right panel, includes legacies operating on the same continent of the airport where slots are reduced. These carriers show modest and stable negative valuations, lacking the initial hub-related advantages but still benefiting slightly from reduced competition and rising fares. In the bottom-left panel, legacies on the other continent show small, slightly negative valuations, reflecting limited ability to leverage market power abroad. American Airlines and Delta exhibit a temporary increase in valuations under moderate constraints, reflecting their greater ability to absorb shocks at an airport within a multi-hub network. Finally, the bottom-right panel, shows that low-cost carriers operating on the same continent consistently gain from slot restrictions by attracting passengers from legacies with cheaper direct flights. Their valuations are consistently negative, i.e. they benefit from less slots at their competitors hub airport.

5.2. *Local environmental scenarios*

We now consider the perspective of regional regulators, reflecting how slot restrictions often stem from local initiatives. The goal is to assess how a locally motivated authority may weigh consumer surplus, airline profits and environmental and health impacts within its jurisdiction. We distinguish between two scenarios. In the first scenario, the regional policy maker focuses only on the costs and benefits to their local residents. Noise and local air pollution are taken into account as they directly impact the surrounding communities whereas CO_2 emissions have mostly non-local effects hence are not included in the welfare calculation. Consumer surplus is split between outbound and inbound flights, with only the share related to the local airport included as relevant consumer surplus. Profits, similarly, are restricted to those generated by locally-based carriers operating from the airport. For airlines with operations spanning multiple hubs, profits are assigned to local regulators according to the proportion of flights conducted at their

[16] From the sensitivity analysis in Appendix F it is clear that the impact of reduced passenger volumes dominates the market power effect and airlines lose from slot restrictions beyond 50%.

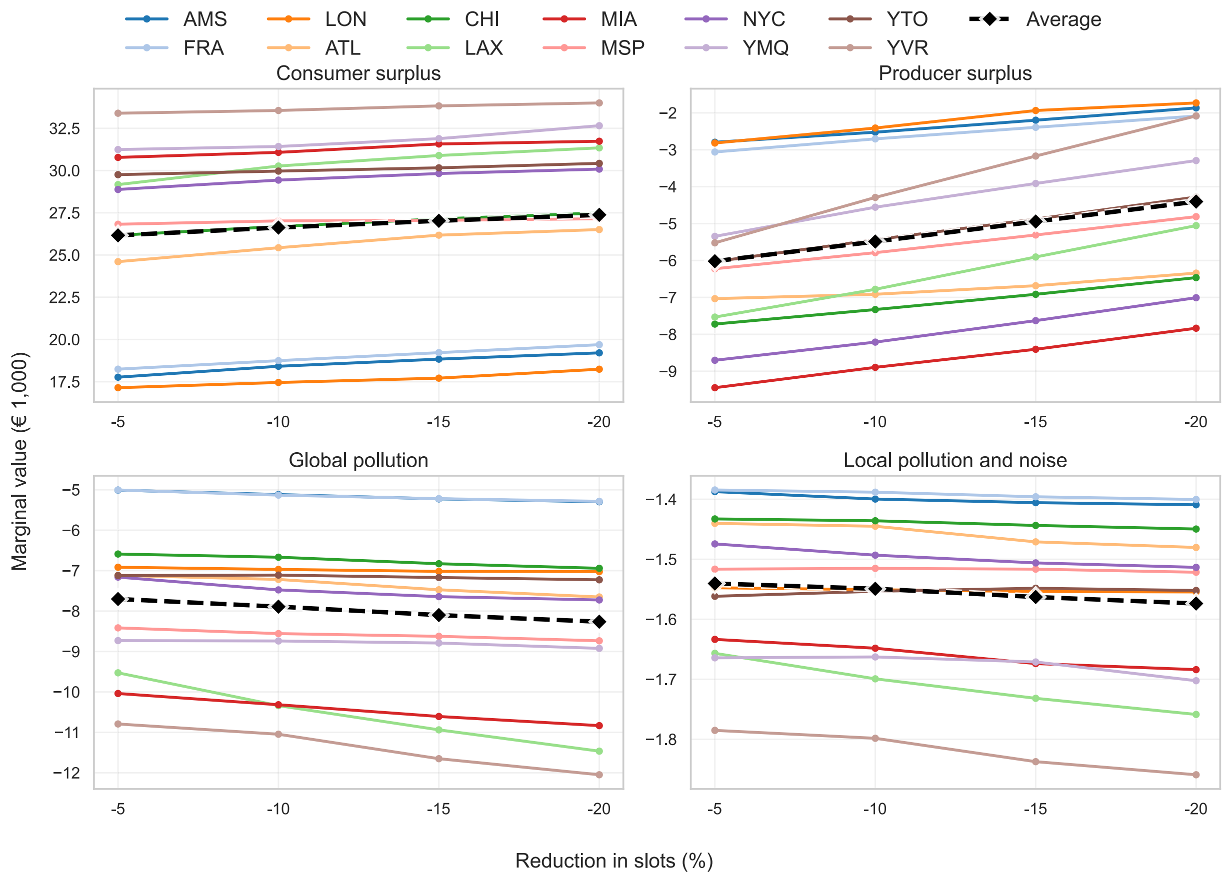

Fig. 6. Components of marginal welfare as a function of slot reductions.

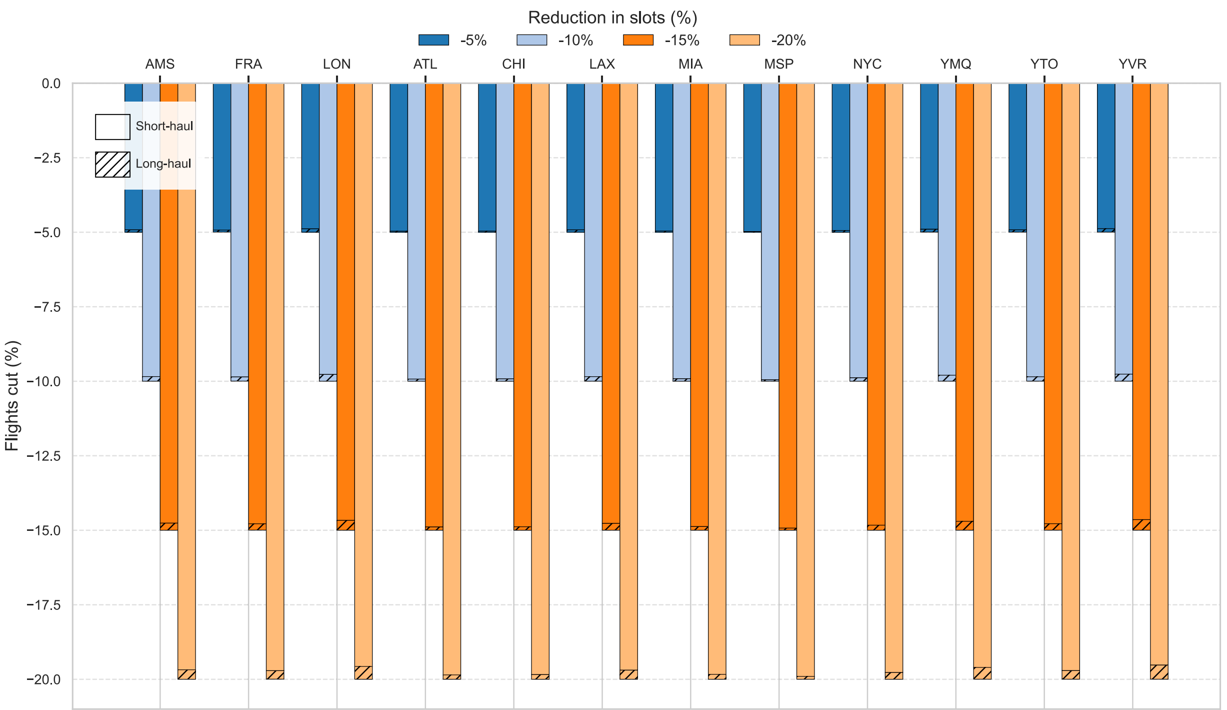

Fig. 7. Flights cut between short and long-haul operations at major hubs.

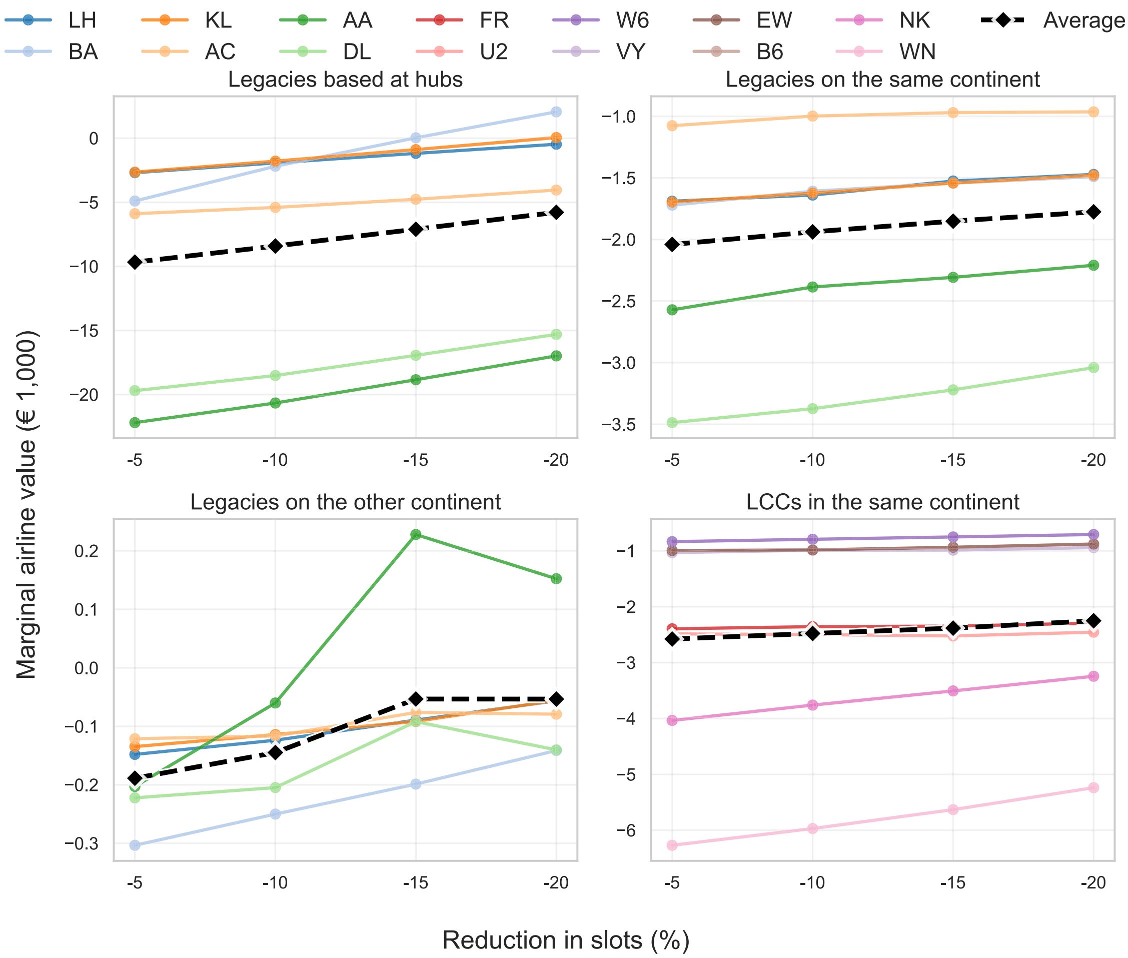

Fig. 8. Marginal slot value for airlines as a function of airport type.

respective hubs relative to their total flight operations. This scenario is formalized by adapting Eq. (21), as follows:

$$w_{\hat{\imath}} = \sum_{\hat{\imath}j\in\mathcal{N}}\sum_{t\in\mathcal{T}}\frac{d_{\hat{\imath}j}}{-\beta_{price}}\ln\left(1+\left(\sum_{a'\in\mathcal{A}}\exp\left(\frac{V_{\hat{\imath}jta'}}{1-\rho_t}\right)\right)^{1-\rho_t}\right)+\sum_{\hat{a}\in\hat{\mathcal{A}}}\pi_{\hat{a}}\left(f^*_{k\hat{a}},p^*_{ijt\hat{a}},\tilde{\alpha}^*_{h\hat{a}},s^*_{k\hat{a}}\right)-\sum_{a'\in\mathcal{A}}\sum_{\upsilon\in\Upsilon}\sum_{\lambda\in\Lambda}\epsilon_{\upsilon\lambda}\xi_\lambda f^*_{\upsilon a'} - \sum_{a'\in\mathcal{A}}\sum_{\upsilon\in\Upsilon}\eta_\upsilon f^*_{\upsilon a'} \tag{22}$$

where $\hat{\imath}\in\hat{\mathcal{N}}$ denotes the slot-restricted node; $\hat{a}\in\hat{\mathcal{A}}$ represents the airlines based at the constrained hub; and $\upsilon\in\Upsilon$ denotes the flight legs connected to node $\hat{\imath}$.

In the second scenario, the regional policy maker aligns with global climate agreements and considers the full scope of environmental impacts associated with slot reductions. This includes not only the environmental impact of flights at the constrained hub but also network effects triggered by reducing slots at that hub, recognizing that CO_2 and other persistent pollutants contribute to climate change regardless of where they are emitted. This is formalized in Eq. (23), as follows:

$$w_{\hat{\imath}} = \sum_{\hat{\imath}j\in\mathcal{N}}\sum_{t\in\mathcal{T}}\frac{d_{\hat{\imath}j}}{-\beta_{price}}\ln\left(1+\left(\sum_{a'\in\mathcal{A}}\exp\left(\frac{V_{\hat{\imath}jta'}}{1-\rho_t}\right)\right)^{1-\rho_t}\right)+\sum_{\hat{a}\in\hat{\mathcal{A}}}\pi_{\hat{a}}\left(f^*_{k\hat{a}},p^*_{ijt\hat{a}},\tilde{\alpha}^*_{h\hat{a}},s^*_{k\hat{a}}\right)-\sum_{a'\in\mathcal{A}}\sum_{\upsilon\in\Upsilon}\epsilon_\upsilon\zeta f^*_{\upsilon a'} - \sum_{a'\in\mathcal{A}}\sum_{\upsilon\in\Upsilon}\sum_{\lambda\in\Lambda}\epsilon_{\upsilon\lambda}\xi_\lambda f^*_{\upsilon a'} - \sum_{a'\in\mathcal{A}}\sum_{\upsilon\in\Upsilon}\eta_\upsilon f^*_{\upsilon a'} \tag{23}$$

As in Eq. (21), the regulator in this setting accounts for all environmental components, including global CO_2 emissions however, consumer and producer surplus are computed for the local population ($\hat{\imath}\in\mathcal{N}$) and the airlines located in the jurisdiction ($\hat{a}\in\mathcal{A}$).

The two scenarios, with results presented in Fig. 9, represent lower and upper bounds on the accountability for environmental externalities by the regional regulator. In both cases, slot constraints at major hubs reduce local social welfare, yielding positive and rising marginal slot values similar to the baseline results. In the first scenario, average marginal slot values range from 11,740 to

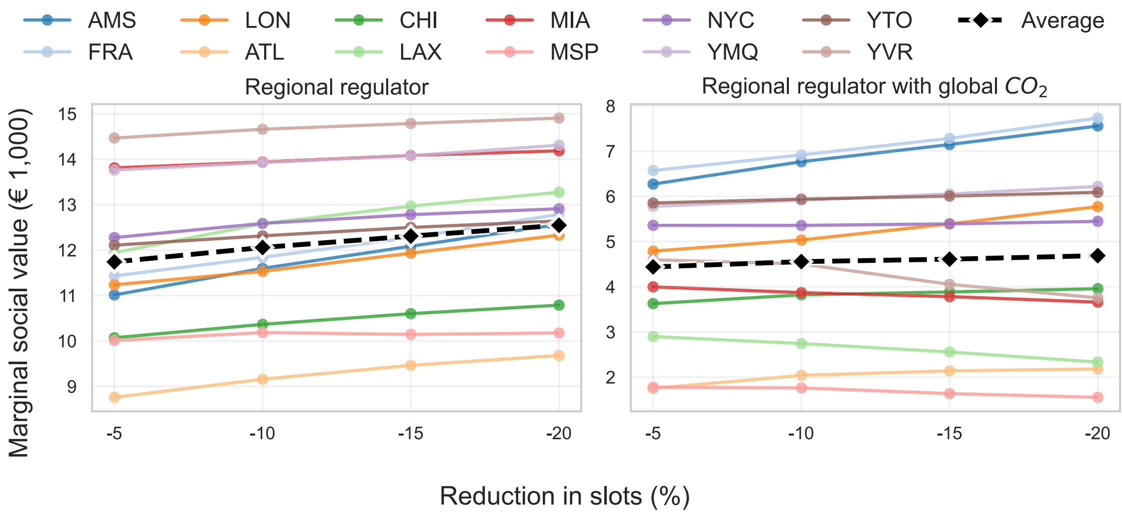

Fig. 9. Marginal value of an additional slot at major airports for a regional regulator.

12,550 euros, with lower valuations at central, dense North American hubs and higher valuations at more remote and less congested airports. European hub valuations fall between these extremes. In the second scenario, reductions in global emissions contribute positively to welfare, particularly when slot constraints are imposed at hubs dominated by long-haul operations such as Los Angeles, Vancouver, Miami and Minneapolis. For these hubs, carbon savings produce a downward-sloping marginal slot valuation. However, tightening slot limits at central hubs still yields negative welfare effects, as losses in consumer surplus outweigh environmental gains. For the average hub in the network, abatement gains bring marginal slot values close to zero at several airports, indicating some policy effectiveness, but remain insufficient to generate positive net welfare.

5.3. Higher social cost of carbon

We apply a social cost of carbon of €200 per tonne, which aligns with common estimates found in the academic literature (e.g., Pörtner et al., 2022). However, considerable uncertainty exists regarding the precise value and many studies propose higher figures for the carbon cost in order to reflect the potentially serious, long-term consequences of CO_2 emissions (Moore et al., 2024; Pindyck, 2019; Rennert et al., 2022). Therefore, we investigate the welfare effects of a slot reduction policy using a carbon cost (ζ) of €400 per tonne. This value is based on the range suggested by Rennert et al. (2022), allowing us to test whether a higher social cost of carbon changes the welfare outcomes discussed so far.

Fig. 10 shows the results of this sensitivity analysis. Higher carbon costs enable airports with a large share of long-haul flights, such as London and Miami, to achieve social welfare gains even with modest slot reductions. These hubs combine long-distance routes with high environmental impacts and strategic positions that generate substantial connecting traffic with relatively low consumer surplus, creating scope for net benefits from slot reductions. Peripheral hubs such as Vancouver and Toronto continue to exhibit substantial declines in social welfare even under a higher carbon cost. Since long-haul services at these airports are reduced earlier when slots are restricted, they realize larger carbon savings, however these gains are more than offset by losses in consumer surplus, reflecting their predominantly origin-destination traffic.

At €400 per tonne of CO_2, regional regulators that fully internalize the network-wide carbon impact of slot reductions almost always see positive local welfare effects from modest restrictions. The exception is certain European hubs dominated by short-haul operations, where carbon savings do not offset the loss in consumer surplus.

5.4. Network centrality

An analysis of network centrality reveals the relationship between marginal social value and airport location within the network.[17] Each airline's set of node-level centrality values is standardized with zero mean and unit variance for direct comparison and

[17] We compute a weighted closeness centrality for each node $i \in \mathcal{N}$ by identifying all origin-destination pairs (i, j) served by airline a. Let E_{ia} denote the set of such itineraries incident to node i. Each itinerary (i, j) is associated with a distance $dist_{ija}$, defined as the sum of the great-circle segments along the path between i and j for airline a. We then define the weighted closeness centrality of node i for airline a as $C_{ia} = \left(|E_{ia}| - 1\right) / \left(\sum_{(i,j)\in E_{ia}} dist_{ija}\right)$, where subtracting one in the numerator ensures that a node connected only to itself has zero centrality. Thus, C_{ia} increases when node i serves many itineraries relative to the total distance of those itineraries. To compare across nodes, we normalize via a z-score transformation. Let $\mu_a = (1/|N_a|)\sum_{i\in N_a} C_{ia}$ and $\varsigma_a = \sqrt{(1/|N_a|)\sum_{i\in N_a}\left(C_{ia} - \mu_a\right)^2}$ where $N_a \subseteq \mathcal{N}$ is the set of nodes served by airline a. The normalized centrality for node i is then $z_{ia} = (C_{ia} - \mu_a)/\varsigma_a$.

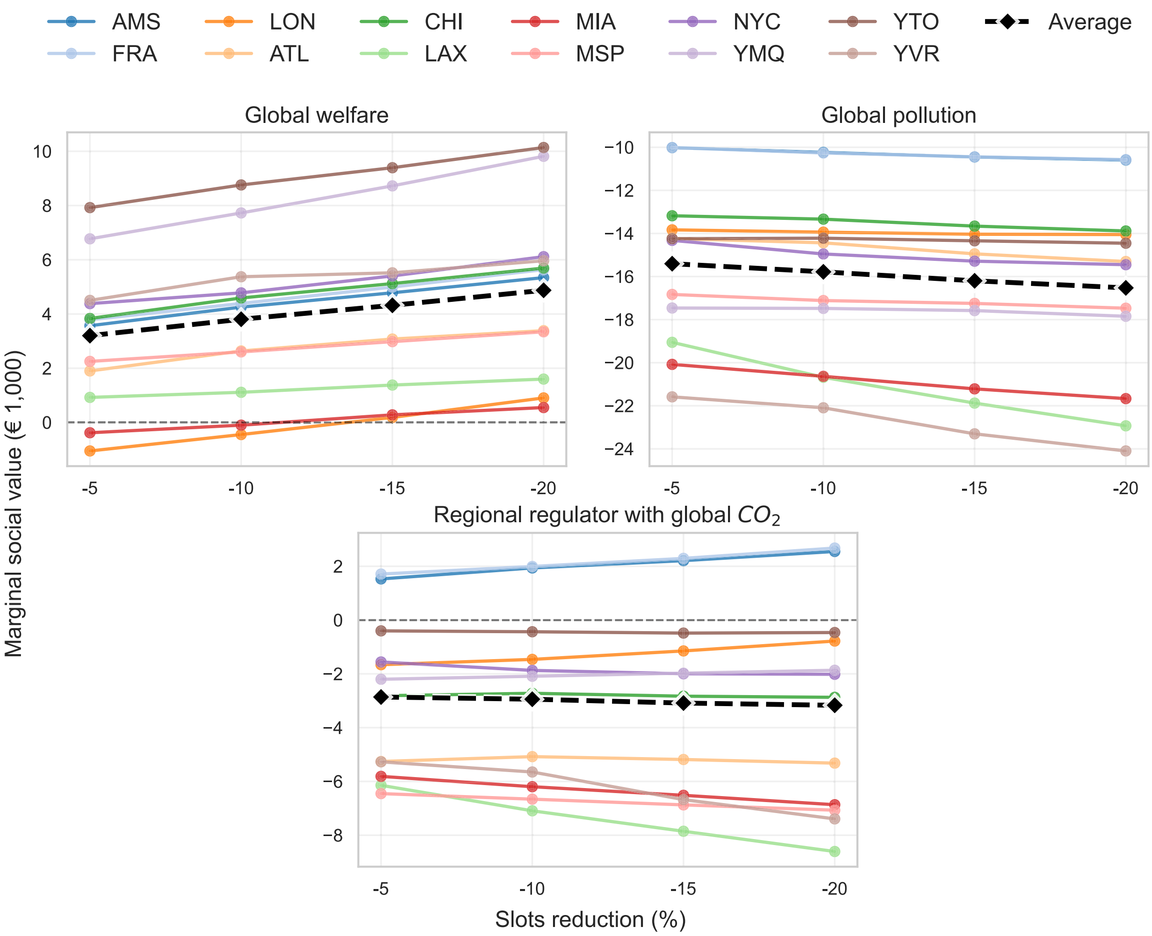

Fig. 10. Marginal social value of an airport slot from the perspective of global welfare, global pollution and of a regional regulator assuming a CO_2 social cost of 400 €/t.

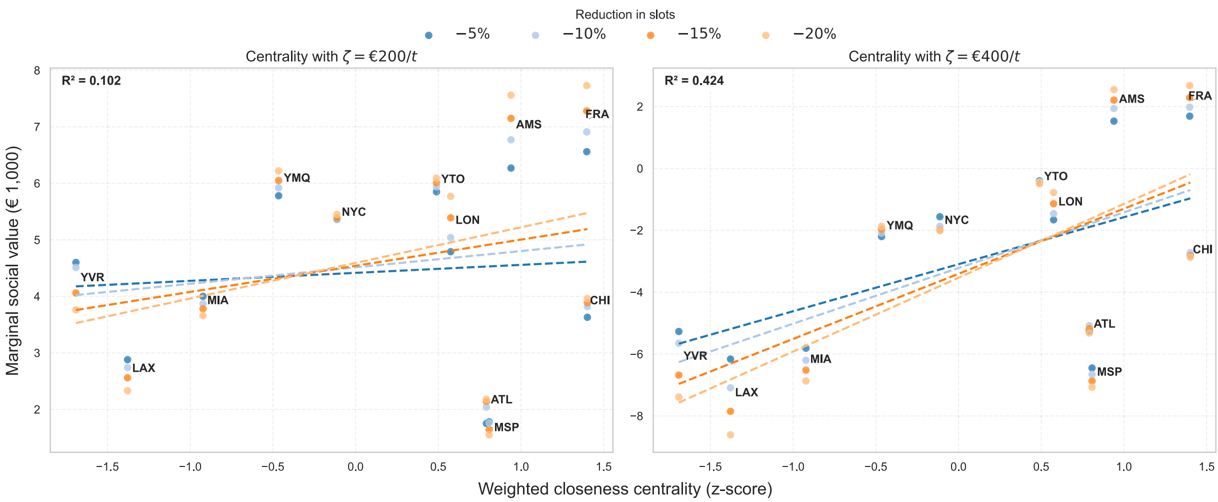

Fig. 11. Marginal social value of a slot as a function of network centrality for different values of the social cost of carbon (ζ).

aggregation into a composite city centrality index. Fig. 11 plots marginal social slot values against the z-score of weighted closeness centrality. Central cities with many short-haul links appear on the right, while peripheral cities with fewer or longer connections lie on the left.

When the social cost of CO_2 is set at $\zeta = 200$, the marginal social value of an additional slot is positive for all constrained hubs and there is no strong relationship between airport centrality and slot value. Environmental savings are insufficient to justify slot restrictions, even at peripheral hubs. When the social cost is higher ($\zeta = 400$), marginal social values become increasingly negative for peripheral cities and a clear linear relationship emerges between centrality and slot value. Nodes with low centrality exhibit the most negative values, while those with high centrality have values closer to zero or lie in the positive domain. This pattern reflects greater emissions savings from reducing longer-distance operations. Removing slots at peripheral nodes, which support longer and less direct flights, yields net societal gains more quickly. Central hubs, handling many short-haul connections, show less negative values due to smaller environmental gains. Under tighter environmental cost assumptions, network structure matters and peripheral nodes become candidates for capacity reduction, while central nodes remain more socially valuable.

6. Conclusions and future directions

We develop a methodological approach for estimating the societal value of airport slot capacity by integrating empirically estimated passenger demand parameters from nested logit models across European, North American and Trans-Atlantic markets into a game-theoretic framework capturing strategic airline competition. Through systematic analysis of slot reduction scenarios ranging from 5% to 20% at major hubs, we quantify the marginal value of airport slots across different stakeholder perspectives, incorporating consumer surplus, airline profits and environmental externalities. Our findings reveal that slot restrictions generate environmental benefits and increase airline profitability by enhancing market power. However, these gains are offset by substantial consumer welfare losses arising from reduced connectivity and higher fares, resulting in net negative social welfare effects at most hubs under realistic carbon valuation assumptions.

Our demand estimations reveal structural differences across markets. We find that price sensitivity differs across regions and segments-North American economy passengers are more price sensitive than Europeans, while the opposite holds in business travel-and that Europeans strongly prefer more direct services. These estimates, derived from a consistent empirical framework applied across six market-segment combinations, may prove useful for further research in these relatively mature markets that represented approximately 50% of revenue passenger kilometres flown in 2024.[18]

Market response analyses show that legacy carriers hubbing at slot-reduced airports benefit from mild slot reductions due to enhanced market power because the additional constraints are applied to all competitors. Conversely, intercontinental carriers based on the other continent are most adversely affected because they directly lose relatively more profitable long-haul routes. Legacy carriers operating within the same continent demonstrate more flexibility, cutting short-haul flights but retaining connectivity and spreading market power gains across multiple routes. Low cost carriers consistently gain from slot restrictions at hubs by attracting passengers with cheaper direct flights, from their legacy competitors that are restricted in their hub-and-spoke network structure.

From a regional regulatory perspective, we distinguish between local and global emissions accountability. When regional regulators consider local costs and benefits alone, consumer surplus losses from reduced connectivity and higher fares consistently outweigh benefits from reduced pollutants and noise. However, when regional regulators also consider global environmental impacts, slot constraints occasionally yield positive welfare effects, suggesting that the level of environmental accountability significantly influences the welfare implications of capacity restriction policies.

To validate our conclusions, we conduct a series of sensitivity analyses on parameter values, competitive structures and alternative control mechanisms, and we explore network centrality as a determinant of slot valuations. The equilibria outcomes prove to be robust to changes in the competitive structure, such as introducing joint ventures that do not alter the core dynamics of the model. Carbon charges are welfare-detrimental because airlines largely pass the additional costs on to passengers. By contrast, hub centrality and assumptions about the marginal social cost of carbon do materially influence the cost-benefit analysis of slot reductions. Under sufficiently high carbon prices, peripheral cities could achieve net welfare gains from additional slot constraints, because the long-haul, more environmentally damaging operations are curtailed more rapidly, generating larger environmental benefits.

The methodology developed in this research, and the accompanying findings, are relevant for policymakers and industry stakeholders, as they highlight the need for comprehensive approaches to slot allocation and management that balance passenger and airline interests with broader societal impacts. The methodology provides policy makers with a tool to understand the distributional effects of slot constraint policies across a network a-priori. Future research could examine policy mechanisms that better align airline incentives with societal benefits and investigate how technological change and evolving travel patterns affect slot valuations over time. It could also explore the interaction between carbon charges and slot constraints to assess how combining quantity and price instruments shapes both slot valuations and welfare outcomes. Furthermore, the model could be extended to incorporate alternative transportation modes for continental destinations, such as high-speed rail, enabling analysis of flight bans for ultra-short domestic routes, as recently proposed in several European countries. Extensions to other regional networks, including Asia-Pacific routes, would enable analysis of a fully global network and address the current limitation that hubs such as Los Angeles and Vancouver appear peripheral despite their centrality to trans-pacific markets.

[18] According to IATA. URL accessed on 30.11.2025: https://www.iata.org/en/pressroom/2025-releases/2025-01-30-01/

In addition, potential future directions could include assessing the impact of new aircraft technologies, such as electric aircraft, hydrogen propulsion and sustainable aviation fuels, on the social value of airport slots, particularly for short-haul routes where these technologies may first achieve commercial viability. The emergence of urban air mobility and autonomous air taxis presents another frontier, requiring analysis of how these new modes interact at slot-constrained airports. From a policy design perspective, research could explore hybrid auction allocation mechanisms that combine market efficiency with social objectives, such as maintaining connectivity to under-served regions while achieving environmental key performance indicators. Finally, behavioural analyses examining how slot constraints differentially affect passenger segments across income levels and geographic regions may provide insights into the distributional consequences of capacity management policies, ensuring that efficiency gains do not come at the expense of equitable access to air transportation.

CRediT authorship contribution statement

Nicole Adler: Writing – review & editing, Validation, Supervision, Methodology, Funding acquisition, Conceptualization; **Gianmarco Andreana:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Conceptualization; **Gerben De Jong:** Writing – review & editing, Writing – original draft, Validation, Software, Methodology, Formal analysis, Data curation, Conceptualization.

Data availability

The authors do not have permission to share data.

Acknowledgements

The authors would like to thank the participants of INFORMS 2024, ITEA 2024 and the 4th SOAR conference for their helpful comments. Nicole Adler and Gianmarco Andreana would like to thank the Israel Science Foundation for grant 2441/21 which helped to finance this research. Nicole Adler would also like to thank the Goldman Center for Data-driven Innovation at Hebrew University Business School for partial funding. Gerben de Jong would like to thank the Netherlands Organisation for Scientific Research (NWO) under the Rubicon programme (award 019.201SG.019) for partial funding.

Appendix A. Handling missing fare data

Missing fare data is a pervasive problem in our dataset and generally for estimating air transport demand models outside the US domestic market. In the OAG dataset used in this paper, 30% to 50% of the fares are missing depending on which region and passenger segment is considered. The fares for LCCs in the US domestic market have the most missing information. To address this issue, we combine supplementary data from DB1B, product-market averaging and the results of imputation methods.

To begin we use the DB1B data to supplement fares in the US domestic market and average all fare observations at the market-product level. For the remaining missing fares, we check whether a fare is available in the “mirroring market” e.g. if the fare is missing for a flight operated by Vueling from Amsterdam to Barcelona, but fare data is available for the Vueling flight from Barcelona to Amsterdam, we consider those fares to be the same.[19] We further eliminate flight alternatives with less than 10 passengers per quarter, for which fares are often missing and which can be assumed not to be representative product alternatives for the majority of travellers. These steps already resolve a considerable amount of the missing data, leaving 5 to 20% missing fare observations for the three regional markets.

We face the remaining absence of data by considering three methods for handling missing data in the presence of endogeneity (Mcdonough and Millimet, 2017). The first, *ad hoc* approach is to drop the products for which the fare is missing. This approach is not optimal, leads to an efficiency loss due to a smaller sample size and may bias the parameter estimates and additional measures obtained from the model such as price elasticities and market shares.[20]

A second approach is to impute the market-level average fare for missing observations. Let w_{ija} be an indicator equal to one if the fare for product a in market ij is missing and replace the fare observations by:

$$p_{ija} = \begin{cases} p_{ija} & \text{if } w_{ija} = 0 \\ \hat{p}_{ij} & \text{if } w_{ija} = 1 \end{cases} \tag{A.1}$$

where $\hat{p}_{ij}$ is the average fare in market i, j. While there might not be much ground to assume that all missing fares would equal the market average, this approach has the benefit of not having to drop any products from the data.

[19] This way of imputation is supported by historical data. Specifically, in markets where both actual and mirroring fares are observed, the interquartile range (25th to 75th percentile) of the absolute deviation between the two is between 10 and 22.55% of the actual fare, depending on consumer segment. This indicates that mirroring fares closely approximate actual fares for the bulk of the observations.

[20] These steps already take us below the 1% in terms of the share of passengers for which the fares are missing (this holds for all regional markets). However, as products chosen by a low number of passengers do still play a role in the denominator of the logit calculations (i.e., it also provides information), this is not the correct way of dealing with this missing data issue.

The third approach, which we ultimately adopt, is to use regression-based imputations. These methods condition fare averages on a set of observed market and product level attributes that are known to impact fare levels. To this end, we use two linear fare regressions, one for economy and one for business fares, in which the regressors are similar to those in our nested logit model, except that we use an indicator for low-cost airlines instead of airline fixed effects (some airlines have too few observations to include airline fixed effects) and also include the temperature difference between endpoint cities (as a proxy for economy flights) and whether a market is domestic. By replacing missing data with values predicted by these regression models, the resulting dataset covers the entirety of airline products.

Appendix B. Algorithm

This appendix summarizes the best-response algorithm to compute the Nash equilibrium of the game Γ (Algorithm 1). Starting from an initial solution, we iteratively solve the non-linear profit maximization problem per airline until successive solutions differ by less than 1%, and we verify robustness by repeating the procedure for 50 random start points and player orderings, as illustrated in Fig. B.1.

Algorithm 1 Algorithm (find the nash equilibria (NE) of Γ).

Start
$solution_0 \leftarrow [0, \dots, 0]$ of length $|\mathcal{A}|$
initialize airlines' decision variables and networks
for each airline in $\mathcal{A}$ **do**:
    solve the mathematical program using *IPOPT* and append to *solution*
end for
while $|solution - solution_0|$ > threshold (i.e. not a best-response) **do**:
    $solution_0 \leftarrow solution$
    for each airline in $\mathcal{A}$ **do**:
        solve the mathematical program using *IPOPT* and replace value in *solution*
    end for
end while
return *solution*
Stop

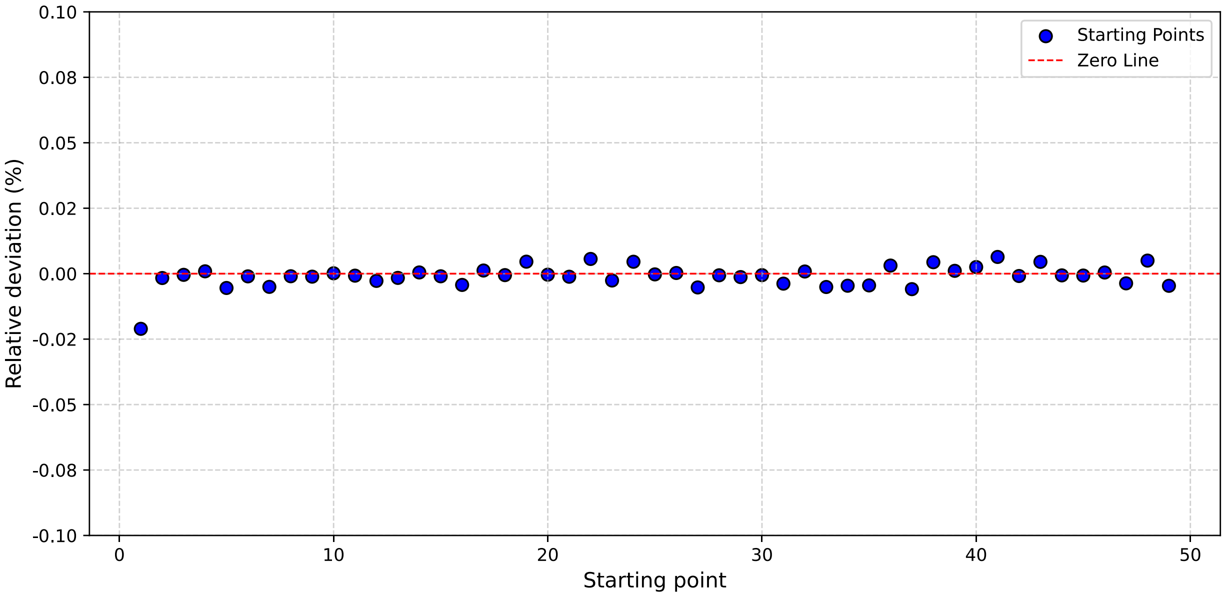

Fig. B.1. Variation of results when initialized from 50 different starting points and player orderings.

Appendix C. Variation of airfares under slot constraints

Fig. C.1 illustrates how slot constraints affect airfares for European and North American carriers. The graphs differentiate between short and long haul flights, further categorized into economy and business segments.

For short haul flights, slot reductions lead to significantly higher airfares, particularly for American low cost carriers. Their market power draws from the non-stop service offered within North America and passengers' higher willingness to pay. On the other hand, European low cost carriers face more price sensitive travellers, limiting their ability to raise fares as easily. Legacy carriers in both regions show similar fare increases.

For long-haul flights, slot constraints raise fares for carriers in both regions and across cabins but particularly in economy class. European carriers increase prices more than their American counterparts because they exploit economies of density at single hub gateways and thus gain market power. In North America the effect of slot reductions is weaker because multi-hub networks mean that passengers are less exposed to constraints at any single airport. Higher fares on both short- and long-haul routes show the additional market power that airlines attain when slot constraints are imposed and their greater ability to control prices under reduced competition.

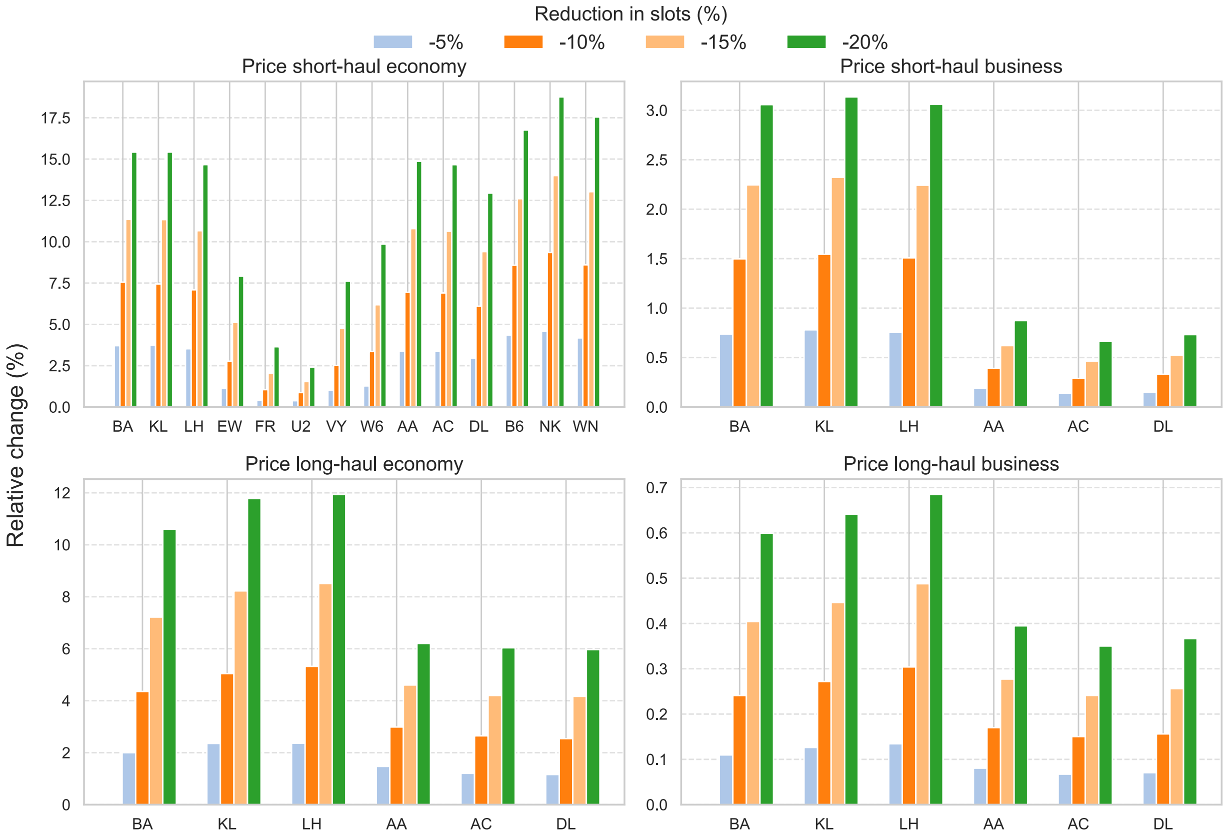

Fig. C.1. Relative changes in airfares for short and long-haul routes among economy and business passengers.

Appendix D. Airline joint ventures

The literature on airline joint ventures shows that network consolidation creates opposing effects. On connecting markets, greater connectivity and the removal of double marginalisation in interline pricing tend to raise consumer surplus (Brueckner, 2001, 2003; Brueckner and Spiller, 1991; Brueckner and Whalen, 2000). By contrast, the market-power effects on overlapping non-stop routes are theoretically ambiguous and empirically mixed (Bilotkach, 2019; Brueckner and Singer, 2019). Reduced competition on some parts of the network may harm consumers and some of the efficiency gains generated by the joint venture are appropriated through higher fares (Adler and Hanany, 2016; Brueckner and Proost, 2010). Overall, some empirical studies find that gains on connecting markets offset the adverse effects on overlapping routes (Brueckner, 2001; Brueckner and Singer, 2019).

Table D.1 compares the baseline with competing legacy carriers to a scenario in which these carriers are consolidated into three joint ventures, each pairing one North American and one European legacy airline. Consolidation softens competition on transatlantic routes. Long-haul frequencies and passenger volumes in both cabins decline by about 10%. At the same time, average long-haul fares increase by roughly 4% in economy and 2% in business, whilst CASK remains unchanged and RASK rises by about one euro cent per ASK. Short-haul markets are less affected with frequencies and passenger number declining by about 4% and average fares increasing modestly in both cabins. These patterns are consistent with a shift in market power on thick long-haul city-pairs, with some spillovers to associated feeder markets within each region, such that total change to airline profits are close to zero.

Table D.1
Comparison of model outcomes between competing legacy carriers (baseline) and joint ventures.

	Competing legacies	Joint ventures
Pax. short econ.	36,298,744	36,533,529
Pax. short bus.	1,492,238	1,514,993
Pax. long econ.	2,664,220	2,528,195
Pax. long bus.	215,442	199,724
CASK (¢)	7	7
RASK (¢)	14	15
Fare short econ. (€)	229	214
Fare short bus. (€)	1949	1942
Fare long econ. (€)	552	565
Fare long bus. (€)	8670	8839
Freq. short	351,067	355,581
Freq. long	9128	8601

Note: “Short” refers to domestic and regional flights within the continent and “long” to transatlantic flights. “econ.” indicates economy class, while “bus.” denotes business class passengers. Passenger numbers and frequencies are aggregated across all airlines, while CASK, RASK and fares represent averages across all airlines weighted by ASK and passengers, respectively.

Fig. D.1 illustrates the distribution of changes in frequencies, and in economy-class passengers and fares under the joint venture scenario relative to the baseline. The distributions reveal heterogeneity across markets. While the consolidation of market power drives price increases in long-haul, hub-to-hub connections, this effect is mitigated by the supply-side efficiencies of the joint ventures in short-haul, hub-to-spoke markets. The higher service frequency offered by the joint venture, relative to individual competing carriers, generates economies of density that stimulate demand in a subset of markets. Consequently, these countervailing anti- and pro-competitive forces tend to offset one another such that aggregate outcomes remain similar to the baseline.

However, our model does not capture interlining synergies, as passengers are not permitted to combine tickets from different airlines to complete their origin–destination journey. Allowing for such itineraries would greatly expand the choice set and render the computation intractable, placing it beyond the scope of this research. Accordingly, the model likely understates the degree of pro-competitive behaviour that would emerge if interlining options were permitted.

Appendix E. Carbon charges

We extend the model to include a carbon charge, ϕ, per tonne of CO_2 produced. The airline profit function is modified accordingly:

$$\max_{\substack{p_{ijta}, f_{ka}, \tilde{\alpha}_{ha},\\ s_{ka}, \hat{m}_{ijsa}}} \tilde{\pi}_a = \sum_{\substack{ij \in \mathcal{N}\\ i \neq j}} \sum_{t \in \mathcal{T}} d_{ij} \hat{m}_{ijta} p_{ijta} - \sum_{k \in \mathcal{K}} c_k f_{ka} - \sum_{h \in \mathcal{H}} o_h \tilde{\alpha}_{ha} - \sum_{k \in \mathcal{K}} \varepsilon_k \phi f_{ka} \tag{E.1}$$

We consider a global scheme that applies the same rate to all flights. Tax revenues enter the welfare function with a positive sign:

$$\tilde{\omega} = \sum_{ij \in \mathcal{N}} \sum_{t \in \mathcal{T}} \frac{d_{ij}}{-\beta_{price}} \ln \left(1 + \left(\sum_{a' \in \mathcal{A}} \exp \left(\frac{V_{ijta'}}{1 - \rho_t} \right) \right)^{1-\rho_t} \right) + \sum_{a' \in \mathcal{A}} \pi_{a'} \left(f^*_{ka'}, p^*_{ijta'}, \tilde{\alpha}^*_{ha'}, s^*_{ka'} \right) - \sum_{a' \in \mathcal{A}} \sum_{k \in \mathcal{K}} \varepsilon_k \zeta f^*_{ka'}$$
$$- \sum_{a' \in \mathcal{A}} \sum_{k \in \mathcal{K}} \sum_{\lambda \in \Lambda} \epsilon_{k\lambda} \xi_\lambda f^*_{ka'} - \sum_{a' \in \mathcal{A}} \sum_{k \in \mathcal{K}} \eta_k f^*_{ka'} + \sum_{k \in \mathcal{K}} \varepsilon_k \phi f^*_{ka'} \tag{E.2}$$

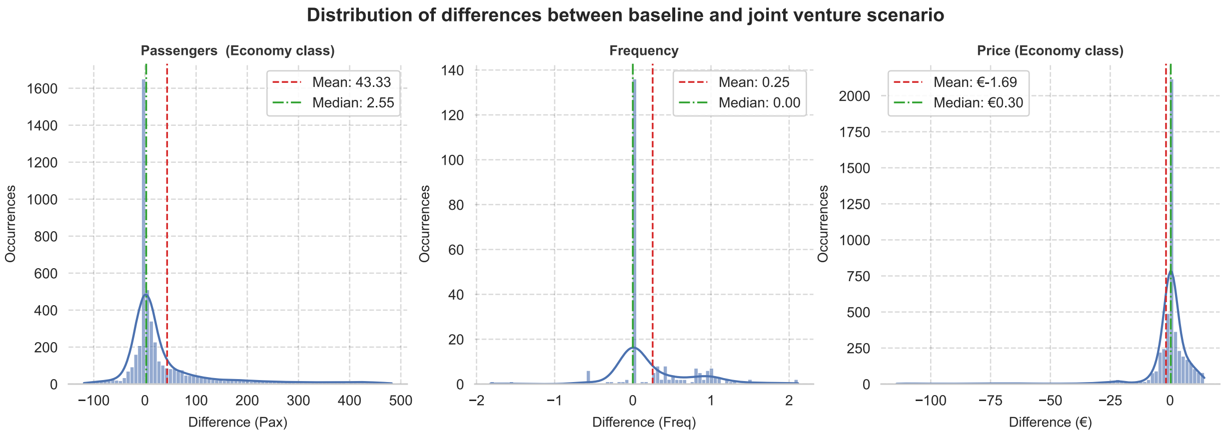

Fig. D.1. Distribution of variations in prices, frequencies, and passenger volumes for the joint venture scenario relative to the baseline.

Table E.1
Welfare impacts of carbon taxation and fleet replacement relative to the baseline.

		Tax 25 €/t		Tax 200 €/t	
	Baseline (€M)	Standard fleet	Upgraded fleet	Standard fleet	Upgraded fleet
Consumer surplus	24,895	−5%	−5%	−38%	−31%
Producer surplus	10,726	−9%	−10%	−50%	−43%
Global pollution	−7,173	−7%	−30%	−50%	−56%
Local pollution	−876	−7%	−35%	−45%	−57%
Noise	−387	−7%	−35%	−50%	−59%
Welfare	**27,185**	**−5%**	**1%**	**−38%**	**−27%**

Note: Welfare values include carbon tax revenues accrued as regulator's surplus. Tax revenues for the standard and upgraded fleets are €88 million and €66 million, respectively, under a 25 €/t tax; these values rise to €379 million and €334 million under a 200 €/t tax.

We apply two carbon price levels, € 25 per tonne (the average EU ETS price in 2019) and € 200 per tonne (a Pigouvian benchmark), assuming a standard fleet and upgraded, more fuel-efficient versions. We assume that a modernized fleet is 25% more fuel-efficient and 30% quieter at the expense of a 10% increase in ownership costs (Axon Aviation, 2022; Easa, 2022; Schäfer et al., 2016).

Table E.1 shows that, absent fleet modernization, welfare declines at both tax levels, with larger losses under the € 200 scheme. The main driver is tax pass-through which encourages airlines with market power to further raise fares, leading to inefficient reductions in demand and a resultant loss in consumer surplus. Furthermore, the reduction in environmental externalities is not sufficient to cover the losses in consumer and producer surplus. Under the assumption of fleet modernization, we see that the low carbon scenario does achieve sufficient environmental externality reductions to lead to an increase in overall social welfare of just over 1%, as compared to that of the higher tax scenario. Unfortunately, whilst a low carbon tax together with fleet modernization may deliver welfare gains, airlines are unlikely to undertake such upgrades voluntarily, given the negative impact on their profits.

Appendix F. Stricter slot constraint scenarios

Fig. F.1 reports the cumulative welfare effects of more extreme slot reductions, with capacity cuts of 50% and 70%. The consumer component declines approximately linearly as slot reductions become stricter, while the environmental components (global and local pollution and noise) improve in a similarly linear fashion. Producer surplus behaves differently. Profits first increase for moderate slot cuts, reflecting stronger market power when all carriers face the same capacity constraint, but then fall once reductions become very large and the loss of traffic dominates. Overall welfare continues to decline with stricter caps because the additional environmental gains do not offset the larger consumer and producer losses.

These scenarios should be interpreted with some caution. Under current operational and regulatory conditions, a reduction in slot capacity of more than 20% at major airports is unlikely. Furthermore, given the airline cost function employed, representing such substantial market changes would require explicitly modelling economies of scope and density, which are not incorporated into our current framework.

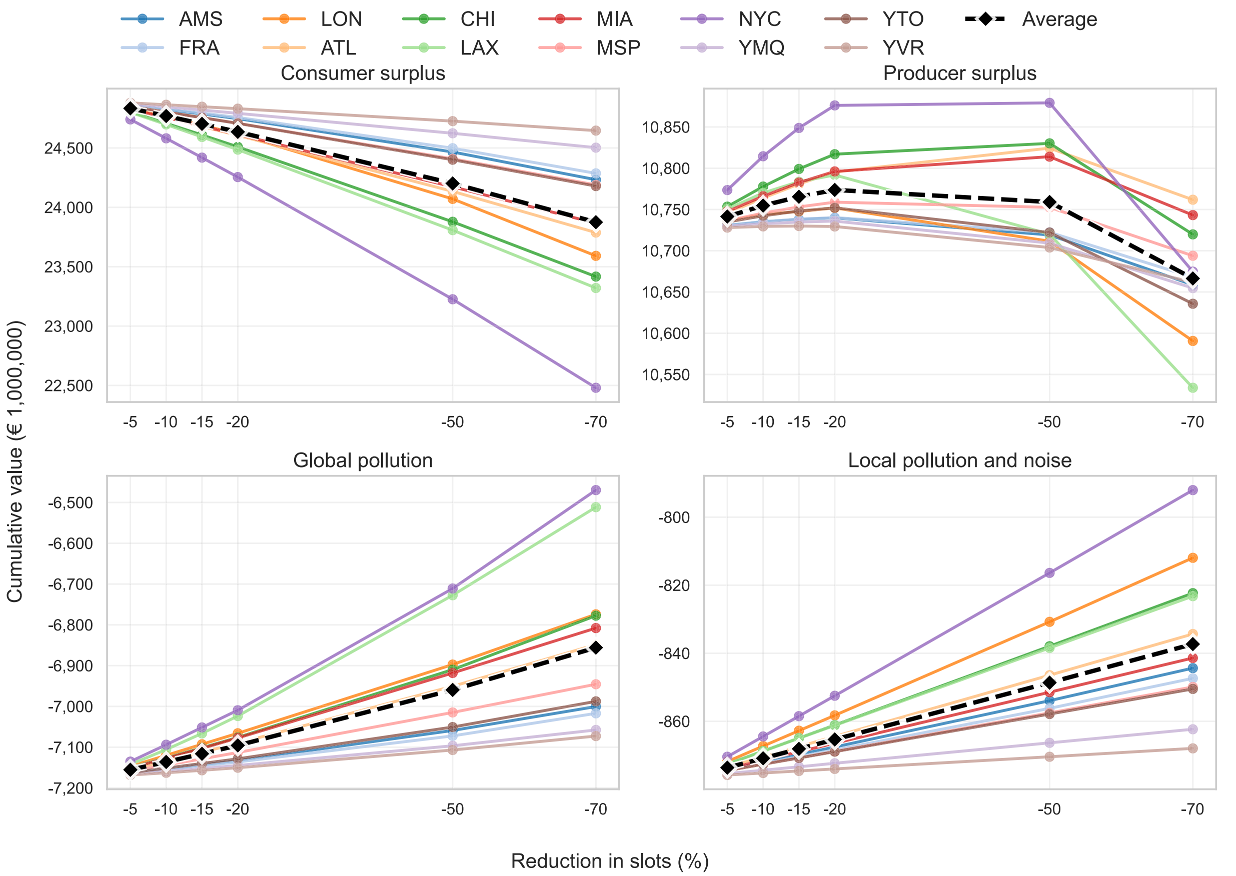

Fig. F.1. Cumulative welfare components for stricter slot reductions.

References

Aviation Week Network, 2016. Oman Air Agrees Record Slot Deal to Grow London Operation. https://aviationweek.com/air-transport/airports-networks/oman-air-agrees-record-slot-deal-grow-london-operation.
Axon Aviation, , 2022. Aircraft Pricing. Accessed: 2023-11-03. http://www.axonaviation.com/commercial-aircraft/aircraft-data/aircraft-pricing/.
Adler, N., 2005. Hub-spoke network choice under competition with an application to western europe. Transp. Sci. 39, 58–72.
Adler, N., Brudner, A., Proost, S., 2021. A review of transport market modeling using game-theoretic principles. Eur. J. Oper. Res. 291, 808–829.
Adler, N., Hanany, E., 2016. Regulating inter-firm agreements: the case of airline codesharing in parallel networks. Transp. Res. Part B Methodol. 84, 31–54.
Adler, N., Martini, G., Volta, N., 2013. Measuring the environmental efficiency of the global aviation fleet. Transp. Res. Part B Methodol. 53, 82–100.
Adler, N., Pels, E., Nash, C., 2010. High-speed rail and air transport competition: game engineering as tool for cost-benefit analysis. Transp. Res. Part B Methodol. 44, 812–833.
Adler, N., Yazhemsky, E., 2018. The value of a marginal change in capacity at congested airports. Transp. Res. Part A Policy Pract. 114, 154–167.
Ball, M.O., Estes, A.S., Hansen, M., Liu, Y., 2020. Quantity-contingent auctions and allocation of airport slots. Transp. Sci. 54, 858–881.
Barnhart, C., Fearing, D., Vaze, V., 2014. Modeling passenger travel and delays in the national air transportation system. Oper. Res. 62, 580–601.
Belobaba, P., Odoni, A., Barnhart, C., 2015. The Global Airline Industry. John Wiley & Sons.
Berry, S., Carnall, M., Spiller, P.T., 2006. Airline hubs: costs, markups and the implications of customer heterogeneity. Competition Policy Antitrust 1, 183–214.
Berry, S., Jia, P., 2010. Tracing the woes: an empirical analysis of the airline industry. Am. Econ. J. Microeconomics 2, 1–43.
Berry, S., Levinsohn, J., Pakes, A., 1995. Automobile prices in market equilibrium. Econometrica 63, 841–890.
Berry, S.T., 1994. Estimating discrete-choice models of product differentiation. Rand J. Econ., 25, 242–262.
Bichler, M., Gritzmann, P., Karaenke, P., Ritter, M., 2023. On airport time slot auctions: a market design complying with the iata scheduling guidelines. Transp. Sci. 57, 27–51.
Bilotkach, V., 2019. Airline partnerships, antitrust immunity, and joint ventures: what we know and what i think we would like to know. Rev. Ind. Organ. 54, 37–60.
Birolini, S., Cattaneo, M., Malighetti, P., Morlotti, C., 2020. Integrated origin-based demand modeling for air transportation. Transp. Res. Part E Logist. Transp. Rev. 142, 102050.
Birolini, S., Jacquillat, A., Cattaneo, M., Antunes, A.P., 2021. Airline network planning: mixed-integer non-convex optimization with demand-supply interactions. Transp. Res. Part B Methodol. 154, 100–124.
Birolini, S., Jacquillat, A., Schmedeman, P., Ribeiro, N., 2023. Passenger-centric slot allocation at schedule-coordinated airports. Transp. Sci. 57, 4–26.
Boonekamp, T., Zuidberg, J., Burghouwt, G., 2018. Determinants of air travel demand: the role of low-cost carriers, ethnic links and aviation-dependent employment. Transp. Res. Part A Policy Pract. 112, 18–28.
Borenstein, S., 1988. On the efficiency of competitive markets for operating licenses. Q. J. Econ. 103, 357–385.
Brueckner, J.K., 2001. The economics of international codesharing: an analysis of airline alliances. Int. J. Ind. Organiz. 19, 1475–1498.
Brueckner, J.K., 2003. International airfares in the age of alliances: the effects of codesharing and antitrust immunity. Rev. Econ. Stat. 85, 105–118.
Brueckner, J.K., 2009. Price vs. quantity-based approaches to airport congestion management. J. Public Econ. 93, 681–690.
Brueckner, J.K., Proost, S., 2010. Carve-outs under airline antitrust immunity. Int. J. Ind. Organiz. 28, 657–668.

Brueckner, J.K., Singer, E., 2019. Pricing by international airline alliances: a retrospective study. Econ. Transp. 20, 100139.
Brueckner, J.K., Spiller, P.T., 1991. Competition and mergers in airline networks. Int. J. Ind. Organiz. 9, 323–342.
Brueckner, J.K., Whalen, W.T., 2000. The price effects of international airline alliances. J. Law Econ. 43, 503–546.
Brueckner, J.K., Zhang, A., 2010. Airline emission charges: effects on airfares, service quality, and aircraft design. Transp. Res. Part B Methodol. 44, 960–971.
Cho, W., Windle, R.J., Dresner, M.E., 2017. The impact of operational exposure and value-of-time on customer choice: evidence from the airline industry. Transp. Res. Part A Policy Pract. 103, 455–471.
Ciliberto, F., Williams, J.W., 2010. Limited access to airport facilities and market power in the airline industry. J. Law Econ. 53, 467–495.
De Jong, G., 2022. Emission Pricing and Capital Replacement: Evidence From Aircraft Fleet Renewal. Available at SSRN 4206318.
Easa, 2022. European Aviation Environmental report. Technical Report. Technical Report.
Eea, 2012. Estimating the External Costs of Industrial Air Pollution Trends 2012-2021. Accessed: 2024-07-04. https://www.eea.europa.eu/publications/the-cost-to-health-and-the/technical-note_estimating-the-external-costs/view.
Eea, 2023. Aviation - Annex 1 - Master Emissions Calculator - 2023 - v1. Accessed: 2024-07-04. https://www.eea.europa.eu/publications/emep-eea-guidebook-2023/part-b-sectoral-guidance-chapters/1-energy/1-a-combustion/1-a-3-a-aviation.3/view.
European Commission, 2020. Handbook on the External Costs of Transport - Version 2019 - 1.1. Publications Office. https://doi.org/10.2832/51388
Fageda, X., Flores-Fillol, R., 2025. The environmental challenge in aviation: can airport charges be part of the solution. J. Assoc. Environ. Resour. Econ. 12, 943–982.
Fageda, X., Teixidó, J.J., 2022. Pricing carbon in the aviation sector: evidence from the european emissions trading system. J. Environ. Econ. Manage. 111, 102591.
Forinash, C.V., Koppelman, F.S., 1993. Application and interpretation of nested logit models of intercity mode choice. Transp. Res. Rec. 1413, 98–106.
Fukui, H., 2019. How do slot restrictions affect airfares? New evidence from the US airline industry. Econ. Transp. 17, 51–71.
Gayle, P.G., 2013. On the efficiency of codeshare contracts between airlines: is double marginalization eliminated? Am. Econ. J. Microeconomics 5, 244–273.
Givoni, M., Rietveld, P., 2009. Airline's choice of aircraft size-explanations and implications. Transp. Res. Part A Policy Pract. 43, 500–510.
Granados, N., Gupta, A., Kauffman, R.J., 2012. Online and offline demand and price elasticities: evidence from the air travel industry. Inf. Syst. Res. 23, 164–181.
Hansen, M., 1990. Airline competition in a hub-dominated environment: an application of noncooperative game theory. Transp. Res. Part B Methodol. 24, 27–43.
Hansen, M., Liu, Y., 2015. Airline competition and market frequency: a comparison of the s-curve and schedule delay models. Transp. Res. Part B Methodol. 78, 301–317.
Hausman, J., Leonard, G., Zona, J.D., 1994. Competitive analysis with differenciated products. Ann. Econ. Stat., 34, 159–180.
Hess, S., Adler, T., Polak, J.W., 2007. Modelling airport and airline choice behaviour with the use of stated preference survey data. Transp. Res. Part E Logist. Transp. Rev. 43, 221–233.
Hong, S., Harker, P.T., 1992. Air traffic network equilibrium: toward frequency, price and slot priority analysis. Transp. Res. Part B Methodol. 26, 307–323.
Iata, 2019. Economic Performance of the Airline Industry. Accessed: 2024-11-22. https://www.iata.org/en/iata-repository/publications/economic-reports/airline-industry-economic-performance.
Iata, 2023. Worldwide Airport Slot Guidelines (WASG). Accessed: 2024-01-25. https://www.iata.org/contentassets/.
Lee, D.S., Fahey, D.W., Skowron, A., Allen, M.R., Burkhardt, U., Chen, Q., Doherty, S.J., Freeman, S., Forster, P.M., Fuglestvedt, J., et al., 2021. The contribution of global aviation to anthropogenic climate forcing for 2000 to 2018. Atmos. Environ. 244, 117834.
Mcdonough, I.K., Millimet, D.L., 2017. Missing data, imputation, and endogeneity. J. Econom. 199, 141–155.
Moore, F.C., Drupp, M.A., Rising, J., Dietz, S., Rudik, I., Wagner, G., 2024. Synthesis of evidence yields high social cost of carbon due to structural model variation and uncertainties. Proc. Natl. Acad. Sci. 121, 2410733121.
Morlotti, C., Cattaneo, M., Malighetti, P., Redondi, R., 2017. Multi-dimensional price elasticity for leisure and business destinations in the low-cost air transport market: evidence from easyjet. Tourism Manage. 61, 23–34.
Mumbower, S., Garrow, L.A., Higgins, M.J., 2014. Estimating flight-level price elasticities using online airline data: a first step toward integrating pricing, demand, and revenue optimization. Transp. Res. Part A Policy Pract. 66, 196–212.
Nevo, A., 2000. Mergers with differentiated products: the case of the ready-to-eat cereal industry. Rand J. Econ., 31, 395–421.
Oum, T.H., Waters, W.G., Yong, J.S., 1992. Concepts of price elasticities of transport demand and recent empirical estimates: an interpretative survey. J. Transp. Econ. Policy, 26, 139–154.
Pindyck, R.S., 2019. The social cost of carbon revisited. J. Environ. Econ. Manage. 94, 140–160.
Pörtner, H.O., Roberts, D.C., Adams, H., Adler, C., Aldunce, P., Ali, E., Begum, R.A., Betts, R., Kerr, R.B., Biesbroek, R., et al., 2022. Climate Change 2022: Impacts, Adaptation and Vulnerability. Technical Report. IPCC Sixth Assessment Report.
Proussaloglou, K., Koppelman, F.S., 1999. The choice of air carrier, flight, and fare class. J. Air Transp. Manage. 5, 193–201.
Pulvino, T.C., 1998. Do asset fire sales exist? An empirical investigation of commercial aircraft transactions. J. Finance 53, 939–978.
Rassenti, S.J., Smith, V.L., Bulfin, R.L., 1982. A combinatorial auction mechanism for airport time slot allocation. Bell J. Econ., 13, 402–417.
Rennert, K., Errickson, F., Prest, B.C., Rennels, L., Newell, R.G., Pizer, W., Kingdon, C., Wingenroth, J., Cooke, R., Parthum, B., et al., 2022. Comprehensive evidence implies a higher social cost of CO2. Nature 610, 687–692.
Reuters, 2023. Dutch Government Scraps Plan to Cap Flights at Schiphol Next Year. Accessed: 2024-06-17. https://www.reuters.com/business/aerospace-defense/dutch-government-suspends-plan-cut-number-flights-schiphol-2023-11-14.
Schäfer, A.W., Evans, A.D., Reynolds, T.G., Dray, L., 2016. Costs of mitigating CO2 emissions from passenger aircraft. Nat. Clim. Chang. 6, 412–417.
Silva, H.E., Verhoef, E.T., 2013. Optimal pricing of flights and passengers at congested airports and the efficiency of atomistic charges. J. Public Econ. 106, 1–13.
Small, K.A., Rosen, H.S., 1981. Applied welfare economics with discrete choice models. Econometrica J. Econ. Soc., 49, 105–130.
Socioeconomic Data and Applications Center (SEDAC), 2018. Center for International Earth Science Information Network - CIESIN - Columbia University. Gridded Population of the World, Version 4 (GPWv4): Population Count, Revision 11. 2018. Palisades, New York: NASA. https://doi.org/10.7927/H4JW8BX5 Accessed: 13-11-2023.
Swan, W.M., Adler, N., 2006. Aircraft trip cost parameters: a function of stage length and seat capacity. Transp. Res. Part E Logist. Transp. Rev. 42, 105–115.
Swaroop, P., Zou, B., Ball, M.O., Hansen, M., 2012. Do more us airports need slot controls? A welfare based approach to determine slot levels. Transp. Res. Part B Methodol. 46, 1239–1259.
Sweeney, R.E., Ulveling, E.F., 1972. A transformation for simplifying the interpretation of coefficients of binary variables in regression analysis. Am. Stat. 26, 30–32.
Vaze, V., Barnhart, C., 2012. Modeling airline frequency competition for airport congestion mitigation. Transp. Sci. 46, 512–535.
Verhoef, E.T., 2010. Congestion pricing, slot sales and slot trading in aviation. Transp. Res. Part B Methodol. 44, 320–329.
Wächter, A., Biegler, L.T., 2006. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Math. Program. 106, 25–57.
Wang, C.H., Zhang, W., Dai, Y., Lee, Y.C., 2022. Frequency competition among airlines on coordinated airports network. Eur. J. Oper. Res. 297, 484–495.
Zografos, K.G., Madas, M.A., Androutsopoulos, K.N., 2017. Increasing airport capacity utilisation through optimum slot scheduling: review of current developments and identification of future needs. J. Sched. 20, 3–24.