Competition and Welfare in Airport Slot Allocation

Sebastián D. Bauer*
Stanford University
Graduate School of Business
sebauer@stanford.edu

November 25, 2025
Updated frequently. Click here for most recent version

Abstract

Airlines require landing and takeoff slots in order to serve congested airports. Currently, airlines are allocated slots for free and have to use them at least 80% of the time to keep them. A frequent proposal is to replace this system with auctions and eliminate the minimum usage requirement. To evaluate the welfare impact of the current system, I build and estimate a structural model of airline demand and route choice using EU data. My main counterfactual compares the status quo to an approximation of the outcome under an auction. I find that the minimum usage requirement leads airlines to increase the total number of flights they offer. This behavior exacerbates congestion, but it also increases competition and the number of available varieties to consumers, thereby increasing total welfare. The current system's pro-competitive effect means that total welfare is actually lower under the auction approximation, although this result reverses if the social cost of carbon is high enough. My results highlight that reforms prioritizing allocative efficiency over the current mechanism's pro-competitive properties risk substantial harm to consumers.

JEL Classification: L13, L93, R41
Keywords: Airport Slots, Competition, Consumer Welfare, Use-it-or-lose-it Rules

*I am indebted to Matt Gentzkow, Mike Ostrovsky, Peter Reiss, and Paulo Somaini for their invaluable support and guidance. I also thank the valuable comments and suggestions of Claudia Allende, Lanier Benkard, Jacob Conway, Ignacio Cuesta, Ben Edelman, Liran Einav, Wes Hartmann, Florencia Hnilo, Allan Hsiao, Zane Kashner, Ravi Jagadeesan, John Lazarev, Neale Mahoney, Marleen Marra, Paul Milgrom, Steve Redding, Adam Rosenberg, Greg Rosston, Al Roth, Daniela Saban, Andy Skrzypacz, Shosh Vasserman, Ali Yurukoglu, participants at various Stanford workshops, and seminar participants at USFCA. I also thank Michele Boyce, Mandar Pradhan, John Radovanovic, Hsin-Pei Wu, and attendees to the 156th IATA Slot Conference for their valuable insights on the airline industry's perspectives on airport slot allocation. I gratefully acknowledge financial support from The Europe Center at Stanford University and the George P. Shultz Dissertation Support Fund at SIEPR.

1 Introduction

Congestion is a widespread issue across most major airports. As of 2025, the International Air Transport Association (IATA) classifies 219 airports worldwide as airports where "demand for airport infrastructure significantly exceeds the airport's capacity" (IATA, 2025). Around 50% of all passengers pass through one of these airports during their trip (IATA, 2024). Furthermore, congestion will only worsen in the coming decades, given that air passenger flow is expected to see a 150% increase by 2043 (Boeing, 2024), and that airport expansion is costly, slow, and politically sensitive in most countries.

To operate at these congested airports, airlines need a slot—the right to land or take off at a specific time. The main way slots are allocated worldwide is through a use-it-or-lose-it rule called grandfathering, whereby an airline keeps a slot for the next year if it used it at least 80% of the time during the current year. If the airline uses the slot less than 80% of the time, then an independent coordinator reallocates it for the following year. At a typical congested airport, most slots are used more than 80% of the time (Bichler et al., 2023).

The effects of the current slot allocation system on welfare are ambiguous. The media, government reports, and researchers argue that the use-it-or-lose-it rule incentivizes airlines to overuse their slots by flying small aircraft and/or nearly-empty flights (GAO, 2012; The Economist, 2017; Kösters et al., 2023). This behavior may in turn suppress entry and harm consumers, which has led to repeated calls for the introduction of airport slot auctions (Ball et al., 2007; Bichler et al., 2021a). On the other hand, use-it-or-lose-it rules can reduce market power by inducing firms to increase production (Gale and O'Brien, 2013; Reitzes et al., 2015). For a concentrated industry like aviation (Borenstein and Rose, 1994; Berry and Jia, 2010), the rule's ability to curb market power could be substantial. The net effect of these opposing forces is a critical open question given recent talks on slot reform by the EU, the UK, and Australia.[1]

In this paper, I analyze how airport slot allocation affects airline behavior and welfare in the EU. I first show reduced-form evidence that the use-it-or-lose-it rule distorts airline behavior but also increases competition. I then construct and estimate a dynamic model of passenger demand, competition and slot usage. I use this model to simulate the system-wide effects of both the current allocation system as well as an auction counterfactual. I find that reallocating slots through an auction and eliminating the use-it-or-lose-it rule results in a 1% drop in total welfare and a 12% drop in total flights. This shows that although the rule distorts airline behavior, it also reduces market power by inducing airlines to increase the total number of flights they offer. This effect is largest for routes between congested airports, as a single flight allows an airline to satisfy usage requirements at both the origin and destination simultaneously.

I focus on the EU for two reasons. First, the EU contains approximately half of the world's congested airports, all operating under a common legal framework. These two facts make the EU an ideal setting to study the system-wide effects of both the current airport slot allocation mechanism as well as alternative mechanisms. Second, the EU is currently considering reforming its airport slot

[1] See European Commission (2023); Government of the United Kingdom (2024); Australian Government (2024).

allocation rules. Therefore, this paper can help inform the EU's policy debate.

I first study the effects of the use-it-or-lose-it rule on airline behavior through an event study. To do this, I leverage the fact that 8 European airports introduced the rule between 2015 and 2019, and I compare them to a set of similar European airports that did not introduce it. My event study reveals that slot controls substantially affect airline behavior. First, airlines increase utilization of previously underutilized slots to approximately 80%. They do this by switching flights from overutilized to underutilized slots, so as to keep as many slots as possible. Second, I find that despite slot controls being seen as a barrier to entry, the share of routes with only one airline falls by 2.1 pp after slot controls are imposed. My results show that the use-it-or-lose-it rule distorts airline choices, but may also lead to more competition.

I next develop a dynamic structural model to quantify the system-wide effects of airport slot allocation. Each period has two stages. In the first stage, airlines—endowed with a set of slots—simultaneously choose their flight schedules. In the second stage, they engage in Nash-Bertrand price competition with differentiated products. The scheduling decision is subject to the constraint that an airline's flight volume cannot exceed its slot capacity at either the origin or destination. This two-stage game repeats every period, with current slot usage determining an airline's future slot allocation through a transition matrix. Airlines are therefore forward-looking, balancing immediate stage profits against the option value of their slot portfolio. I estimate my demand model and cost parameters using data for intra-Schengen[2] passenger flights from 2015 to 2019.

I use my model estimates to simulate how airlines behave under the current system and under counterfactual policies. The dynamic problem is a version of the Generalized Assignment Problem and is thus NP-Hard,[3] so I approximate airline behavior with a stylized assumption to simulate the effects of grandfathering. Specifically, I assume that (i) airlines behave as if they were required to use their slots at least 80% at each slot-controlled airport, and that (ii) airlines believe all other airlines also behave this way.[4] With these two assumptions, solving the equilibrium under the baseline involves maximizing static airline profits with a series of minimum usage restrictions for each slot-controlled airport. These two assumptions are a stylized simplification of airline behavior that fall in line with what I observe in my event study, where airlines strategically change their slot usage to keep most of their slots. In addition, I show that my baseline simulations match my reduced form event study estimates for the 8 airports that start having a usage requirement.

In my main counterfactual, I analyze the effects of eliminating the use-it-or-lose-it rule if a social planner simultaneously reallocates airport slots to maximize gross total profits. This is an approximation to the outcome if slots were assigned using a combinatorial auction. Rassenti et al. (1982) show that such an auction maximizes total bidder's surplus, which in my context is the sum of total (current period) profits. I solve this computationally intractable problem using a genetic

[2]The Schengen Area is an open borders area in Europe. Between 2015 and 2019 the Schengen Area included 22 of the 28 EU members, plus Iceland, Liechtenstein, Norway and Switzerland. See Section 5 for more details.

[3]NP-Hard problems are a class of computational problems for which no known efficient algorithm exists to find an optimal solution.

[4]In Appendix F I test an alternative assumption where airlines are only required to keep the slots they keep in the data, finding broadly similar results.

algorithm that evolves allocations toward the profit-maximizing outcome.[5] In essence, the algorithm acts like a sophisticated form of trial-and-error, where it starts with a set of random ways to assign slots, keeps the most profitable ones, and intelligently combines their best features to "breed" a nearly optimal solution. I rerun this algorithm while sequentially excluding one airline at a time, which allows me to compute the auction's VCG payments.[6]

I find that airlines reduce the total number flights they offer by 12%, resulting in a 21% drop in consumer welfare, a 20% increase in gross airline profits, and a 1% drop in total welfare—the sum of consumer welfare and gross airline profits. The reduction in the total number of flights means that airlines serve fewer markets, which allows them to increase their markups by 44% and offer fewer varieties. Flights between congested airports fall by 14%, which shows that the current system has a particularly strong pro-competitive effect on these markets. Finally, the auction generates around $5.3 billion in annual revenue. This means that profits net of auction payments are 23% lower than the baseline, which may help explain airlines' opposition to auctions (IATA, 2024).

The 12% reduction in the total number of flights has a positive effect on the environment. I estimate that the total saved emissions are around 6 million metric tons of CO_2, or 0.16% of annual EU emissions. This means that auctions may result in a gain in total welfare if the social cost of carbon (SCC) is higher than $51 per metric ton of carbon. Still, I find that consumers are better off under the baseline for most of the usual SCC estimates found in the literature (EPA, 2023), even if they were to accrue the gains from lower carbon emissions.

To decompose the pro-competitive effect of the use-it-or-lose-it rule from the auction's efficiency gains, I run a second counterfactual where I eliminate the rule while holding the current slot allocation fixed. This would be equivalent to a system where airlines retain their currently allocated slots regardless of usage. This leads to a 6% increase in airline profits and a 15% decrease in consumer welfare, resulting in a 5% decrease in total welfare. Comparing the two counterfactual simulations shows that the efficiency gains from reallocating slots amount to around 4% of total welfare. Nevertheless, consumers are worse off under the auction allocation. This falls in line with the theoretical literature on essential input auctions, which shows that standard auction mechanisms can lead to lower competition and consumer welfare in the downstream market (Rey and Salant, 2017).

My analysis shows that tying today's slot usage to tomorrow's slot ownership generates significant consumer surplus by inducing airlines to increase the total number of flights they offer, thus reducing market power and increasing product variety. The effect is particularly strong on routes between congested airports, as airlines can satisfy usage requirements on both origin and destination with a single flight. My findings introduce a critical and previously unquantified trade-off for policymakers. My results demonstrate that reforms focused solely on efficiency gains from slot reallocation could inadvertently sacrifice billions in welfare by diminishing this competitive pressure.

Related Literature. My paper relates to the literature on airport slot market design, which

[5] A recent literature has shown that these type of algorithms can yield significant improvements in airport slot allocation (Melder et al., 2023; Katsigiannis and Zografos, 2024).

[6] The VCG payments are due to Vickrey (1961), Clarke (1971) and Groves (1973). In this setting, they involve computing the expected externality that an airline imposes on other airlines by participating in the slot auction and acquiring some slots.

traces its origins to Rassenti et al. (1982). Since then, the literature has studied how to optimally design landing slot allocation under inclement weather (Schummer and Vohra, 2013; Schummer and Abizada, 2017), and how to allocate slots through an auction or trading for one airport at a time in the US (Ausubel et al., 2006; Ball et al., 2007) and Europe (Condorelli, 2007; De Wit and Burghouwt, 2008; Bichler et al., 2021b, 2023; Kösters et al., 2023). In a recent paper, Marra (2024) builds and estimates a model to compute the value of an airport slot for the two main Paris airports based on the downstream profits, and uses that estimation to calculate the welfare gains from achieving an efficient allocation. My paper's setting with multiple slot controlled airports across different metropolitan areas allows me to measure the system-wide effects of both the current slot allocation method as well as that of an auction outcome. This is a key consideration in my setting, given the interlinked nature of slot valuations.

As my paper's main counterfactual is an approximation to an auction, I also speak to the theoretical literature on essential input auctions. Seminal work by Rey and Salant (2017) shows that in these settings an auction may not result in an efficient allocation if input ownership has increasing returns to scale. In addition, essential input auctions may also be inefficient due to downstream competition (Janssen and Karamychev, 2010), downstream deterrence (Bauer, 2025), or if there are not enough bidders in the auction (Kasberger, 2023). My paper contributes to this literature by empirically showing how usage requirements can be used to fix some of the competitive issues in essential input auctions.

My paper also relates to the extensive literature on airline competition. This includes papers that study dynamic pricing strategies (Borenstein and Rose, 1994; Lazarev, 2013; Williams, 2022; Betancourt et al., 2022; Aryal et al., 2024), papers that study various aspects of market entry (Reiss and Spiller, 1989; Berry, 1992; Ciliberto and Tamer, 2009; Ciliberto et al., 2021), and papers that study the effects of airline mergers (Benkard et al., 2010; Berry and Jia, 2010; Bontemps et al., 2022). Within this literature, my paper is particularly related to a subset of papers that studies different aspects of network competition (Borenstein, 1989; Berry, 1990; Berry et al., 1996, 2006; Aguirregabiria and Ho, 2012; Ciliberto and Williams, 2014; Dix and Orzach, 2022; Bontemps et al., 2023; Yuan and Jia-Barwick, 2024). My paper shows how usage requirements for essential inputs in the airline industry can promote or hinder competition throughout the network.

Finally, my paper also contributes to the literature on use-it-or-lose-it rules. These rules can be found across various essential inputs besides airport slots, such as spectrum rights (GAO, 2014), water rights (Ferguson, 2024), land rights (Raz, 2018), and state budgets (Liebman and Mahoney, 2017). The theoretical literature on the topic has shown that these rules may reduce market power by inducing higher production (Gale and O'Brien, 2013; Reitzes et al., 2015). My paper contributes to this literature by showing that this welfare enhancing effect can be substantial, especially when the inputs subject to the rule are geographically linked.

2 Institutional Background

The current system of airport slot allocation in the EU traces its origins to the Regulations EEC-95/93 and EC-793/2004. Both regulations draw heavily from IATA's Worldwide Airport Slot Guidelines (WASG), a list of recommended rules that IATA has been publishing since the early 1970s. These regulate how an airport is declared slot controlled, how slots are allocated and kept, and what constitutes slot misuse. In this section I provide a summary of the most important aspects of this process. For a thorough and up to date review of the slot allocation process in Europe, see Odoni (2021).

The WASG divide all airports into three levels. Level 1 airports are airports where capacity is adequate at all times. Level 2 airports are airports where some times of the day exhibit congestion, but where mutual agreements between airlines can be made to ease problems. Level 3 airports are airports where demand for airport infrastructure significantly exceeds capacity. This last group is the group where airports are *slot coordinated*, meaning that any airline that wants to operate out of one of these airports needs a landing/take-off slot pair. While in theory landing and take-off slots could be allocated separately, they are always allocated together to a specific pair of inbound/outbound flight numbers. Each level 3 airport has an airport coordinator, who is an independent party in charge of allocating the slots each year.

For slot ownership purposes, each year is divided into two seasons: a Summer season, running from the last week of March to the last week of October, and a Winter season for the remainder of the year. While in general airports have the same level for both Summer and Winter seasons, some European airports have different levels for different seasons. For example, Ibiza airport is Level 3 during the Summer season but Level 2 during the Winter season. Within each season, a slot is a specific day of the week/time of the day combination. But the combination can be operated by a different airline during the Summer season and the Winter season.

Each season, the slot allocation process starts with an airport's capacity declaration.[7] This capacity declaration frequently includes the number of aircraft the airport can accommodate at the various potential bottlenecks, such as terminal gates, the runway, the apron, and the number of baggage handlers, among others. These capacity declarations are usually very granular, with specific number of arrivals and departures for each 5 minute interval in a week. In the EU there are also separate capacity declarations for Schengen and non-Schengen flights, as passengers from the latter have to go through border controls.[8] Figure 1 shows the countries that were part of the Schengen Area over this paper's time period.

Based on this capacity declaration, the airport authority, the independent coordinator or a group of airlines handling a majority of the airport traffic can ask for the airport to go up or down a level. In general, airports have gone from not having slot controls to having them, as capacity has historically not kept up with demand. That being said, there are examples of airports that have stopped being

[7] See Ribeiro et al. (2019) for an example of such a capacity declaration for Lisbon Airport.

[8] A Schengen slot can be turned into a non-Schengen slot if the border constraint is slack at a particular time of the day. This happens on a case-by-case basis, and requires permission from the coordinator.

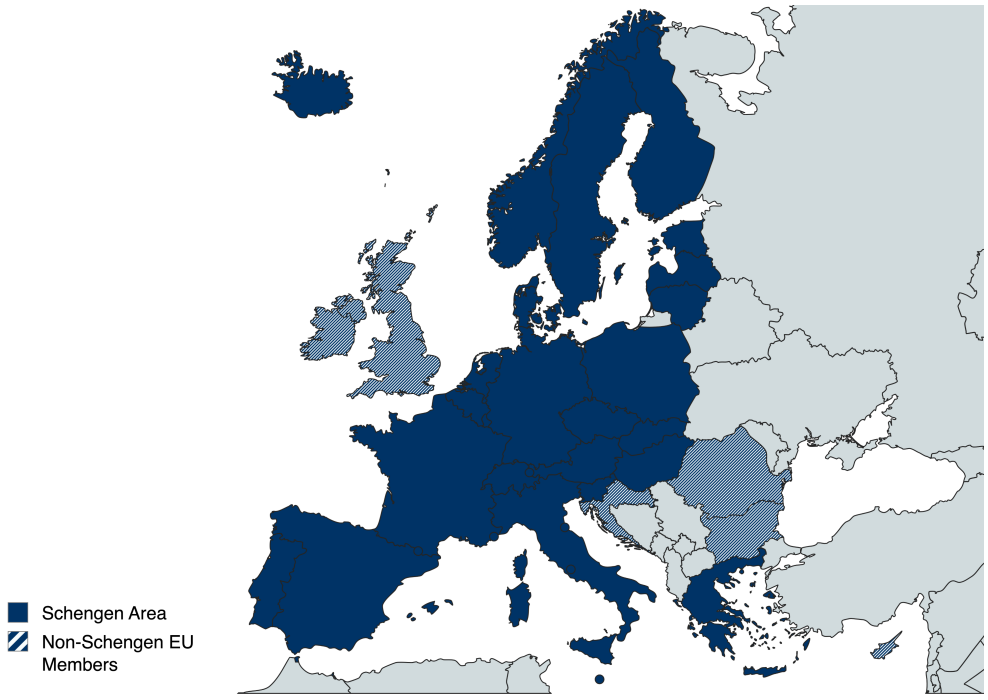


Figure 1: Schengen Area over the 2015-2019 period. The map also includes the 6 EU countries that were not part of the Schengen Area at the time. Since then, Bulgaria, Croatia, and Romania have joined the Schengen Area, whereas the UK has left the EU.

Level 3, such as Newark Airport in 2016 or Lyon Airport in 2021. Whether an airport is then declared to be slot constrained or not is up to the national civil aviation authority. In general, one or more of the relevant stakeholders request the level change when too many conflicts cannot be solved by mutual agreement.

The main allocation happens around two annual meetings, one for each season. During these meetings, coordinators, airline representatives and airport representatives meet to negotiate the slot allocation. Before the meetings, all airlines receive the list of all their grandfathered slots—i.e., the slots they managed to use at least 80% during the last season. In addition, airlines send an initial request to the coordinator with all the new slots they need before the meeting, and some of these requests are granted ahead of the meetings. Thus, the meetings mostly serve to discuss a minority of slots where the airline would like to get an improvement over what it got during the pre-meeting allocation.

While coordinators have a lot of discretion in how to allocate slots, they have to respect two rules. First, they have to respect grandfathering rights: if an airline had a slot last season and used it at least 80% of the time, it is entitled to keep the slot. Second, 50% of all available (non-grandfathered) slots have to give priority to new entrants, which are airlines with less than 5% of slots at the airport. Beyond these rules, there's a list of additional criteria the coordinator may use. This list includes how long the airline has been waiting for a slot, the airline type, the amount of competition and connectivity the flight would bring, as well as local and environmental

considerations. The coordinator has full discretion over which of these criteria to apply for each case.

Airport slots are allocated for free, but airlines have to pay the airport fees each time they operate them. While airlines are generally allowed to exchange slots on a one-for-one basis, within the EU they are not allowed to offer monetary compensation or to lease them. In most cases, these trades are not public information, although the coordinators of some countries do publish them.[9] Even in countries where monetary compensation is allowed, the payment is almost never public information. Thus, there are very few examples of slot sales, all coming from non-EU countries.[10]

An airline has to announce how it will use each slot through a General Clearance Request. This request includes the flight number or the aircraft registration number, the departure/arrival time and airport, and the type of aircraft used. This is what then allows the coordinator to verify that an airline has used a slot more than 80%. While in theory the coordinator should count any flight outside the allocated slot as one instance where the slot was not used, in practice coordinators give airlines some leeway. For example, the coordinator for the Austrian airports allows a 30 minute window for arrival/departure.[11] In addition, airlines are not penalized when they have to cancel flights due to inclement weather or airport-wide problems. The extent to which airlines are granted this Justified Non-Use of Slots (JNUS) varies by coordinator.

3 Reduced-Form Evidence on Usage Rules and Airline Behavior

I take advantage of the fact that between 2015 and 2019 8 Schengen Area airports started having slot controls to show reduced form evidence of the effect that usage rules have on airline behavior. I do this through an event study, where I use other comparable Schengen Area Airports that did not impose slot controls as my control group. I find that slot controls increase the usage in previously underutilized time windows, and they also increase the share of flights to other slot-controlled airports, which leads to a decrease in the share of flights in monopolistic routes. This section's findings serve as motivation for the rest of the paper.

I use data from Eurocontrol's Aviation Data Repository for Research (ADRR) to construct my event study. A detailed explanation of this dataset as well the other data sources that I use throughout the paper can be found in Section 5. In this section I only present the essential information needed to understand the event study. The dataset contains flight level data for all flights that crossed European airspace during the months of March, June, September and December, between 2015 and 2022. The three variables that I study are slot usage, flights to other congested airports, and flights in monopoly routes.

I restrict attention to flights that depart from one of my treatment or control airports, but I consider both Schengen as well as non-Schengen destinations. Table 15 in Appendix I presents the list of treatment and control airports, along with the first season and year when each treatment

[9] For example see Airport Coordination Limited (2025) for a list of completed trades in the UK and the UAE.
[10] See Odoni (2021) for some examples of payments for London-Heathrow slots.
[11] See Schedule Coordination Austria (2024).

airport becomes slot constrained—i.e., when it became a Level 3 airport. My control airports are the 44 airports that were classified as Level 2 during this period. Recall that a Level 2 airport is an airport without slot-controls but with some levels of congestion. In general, Level 3 airports spend some seasons as Level 2 airports before introducing slot controls, which makes these airports a good control group.[12]

I allocate each flight to a five minute time-window on a given day of the week, and I separate flights into two seasons, in line with the IATA seasons: Most of March along with December of the previous year belong to the Winter Season, whereas the last week of March, along with June and September belong to the Summer Season. These 5 minute time-windows during a specific season are the units of observation for my event study. I choose this approach rather than having individual flights as my unit of observation because before slot controls, airlines have much more leeway to modify their schedules, and it is thus harder to construct a series of comparable flights. As both the 8 treatment airports as well as the 44 control airports are medium and small airports, they generally have only one departing flight per 5 minute time window.

I restrict attention to the year immediately before and immediately after slot controls are imposed. As my dataset has only four months of the year and as time windows are defined on a weekly basis, I have 8 observations before slot controls and 8 after for each time window. I restrict my sample in this way because in general, schedules do not vary all that much from one year to the next, but they may vary substantially over longer periods of time. Indeed, across my 8 treatment airports, 94% of all flight series present during the period immediately before slot controls are also present during the period immediately after.

I classify the time windows in bins according to their pre-slot controls average usage—i.e., the share of weeks a time window was used before slot controls are imposed. I do this because imposing slot controls should mainly affect airline behavior in time-windows that were previously underutilized. I place each time window into a 10% bin. This means that the 40% usage bin will include all time windows that were used between 40% and 50% before slot controls are imposed. Thus, for each outcome variable I run several event studies.

Formally, I run the following regression on these time windows, where w represents the time window, and t represents the period (one week):

$$y_{w,t} = \sum_{t'=-8}^{7} \beta_{t'} \mathbb{1}(t' = t - t_0) + \eta_w + \phi_t + \varepsilon_{w,t},$$

where $y_{w,t}$ is my outcome variable of interest, t_0 is the first week of slot controls, and η_w and ϕ_t are time window and week fixed effects, respectively. I cluster my standard errors at the airport level, as that is the level of treatment. As with any staggered treatment event study, my regression may be biased if the treatment homogeneity assumption across airports treated during different periods does not hold. To address this issue, in Appendix C I run all my event studies while applying the

[12]Indeed, between 2020 and 2025 4 further airports in the control group became Level 3. As throughout this paper I only consider pre-2020 data to avoid the Covid-19 years, I omit adding these airports to my treatment group.

robust imputation proposed by Borusyak et al. (2024), finding broadly similar results.

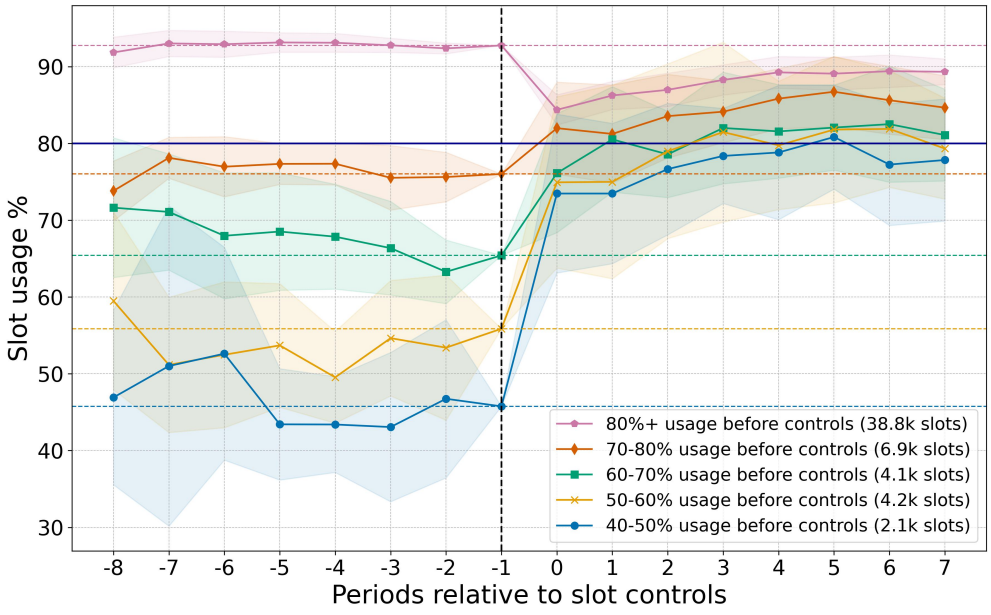


Figure 2: Event studies showing the effect of airport slot controls on time window usage. I compare the 8 Schengen Area airports that started having slot controls between 2015 and 2019 against a comparable set of airports that remained without slot controls during the period. My unit of observation is a 5-minute time window during a specific day of the week during a specific season (e.g. Mondays from 9:00-9:05 during the Summer season). My unit of time is one week. I cluster errors at the airport level. The graphs show all event studies with the baseline added in. The value at -1 (the last period before slot controls) is the baseline value. As I am comparing time windows belonging to the same season, there are 36 weeks between the last period without slot controls and the first period with them. The X-axis advances 2 months every 4 weeks due to the composition of the ADRR dataset.

Figure 2 presents the first of these event studies, where $y_{w,t}$ is equal to a dummy variable that takes value 1 if the time window w was utilized during week t, and 0 otherwise. For this event study, I classify observations into 5 bins: 40-50% usage, 50-60% usage, 60-70% usage, 70-80% usage, and 80%+ usage.[13] To make comparison easier, I add the baseline value to each event study. As I am comparing time windows that belong to the same season during the year immediately before and immediately after the beginning of slot controls, there are approximately 36 weeks between the last pre-treatment period and the first post-treatment period. I find that airlines substantially adjust their slot usage following slot controls. First, airlines increase usage in sub-utilized time windows to approximately 80% usage. They do this by lowering usage on time windows that were previously used more than 80%.[14] This result shows that the grandfathering rule generates substantial distortions in airline behavior.

[13] There are too few observations with less than 40% usage to make meaningful inferences.

[14] As most time windows were being used more than 80% before the change, the overall effect on usage is slightly negative, As can be seen in Figure 20.

Figure 3 presents the impact of slot controls on competition. Here the unit of observation is no longer a time window, but a route. I find that the share of monopoly routes decreases by around 2.1% during the first season with slot controls. Thus, despite slot controls sometimes being seen as a barrier to entry, the fact that airlines increase their share of flights to the same set of airports drives up competition among existing airlines. The estimates are noisier due to the coarser unit of observation, as there are fewer routes than time windows.

Finally, Figure 4 presents the impact of slot controls on the share of flights that go to other slot-controlled airports. We can see that the time windows with usage rates below 75% see a 3.5% increase in the flight share to other slot-constrained airports. We also see that for the group above 75%, there is no significant effect. This result shows that if airlines need to increase their slot usage, they do so mainly to the same set of airports—i.e., those who have slot controls.

These results show that airlines respond to the incentives that slot controls create. Specifically, they change the usage of their slots so as to maintain their slot portfolio. As flying flights between two slot controlled airports protects 2 slots with one flight, this leads them to increase the flight share to other slot-constrained airports. When all airlines do this, the route overlap among airlines increases, which in turn drives up competition.

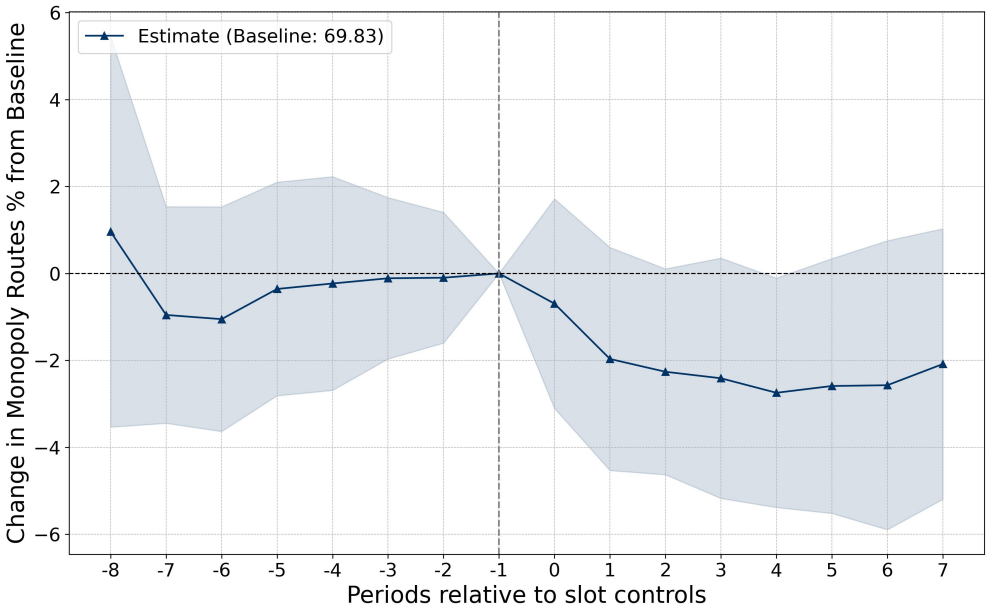


Figure 3: Event studies showing the effect of airport slot controls on the share of routes only one airline serves. I compare the 8 airports that started having slot controls in the Schengen area against a comparable set of airports that remained without slot controls during the period. My unit of observation is a route (origin-destination pair). I cluster errors at the airport level. The value at -36 (the last period before slot controls) is the baseline value.

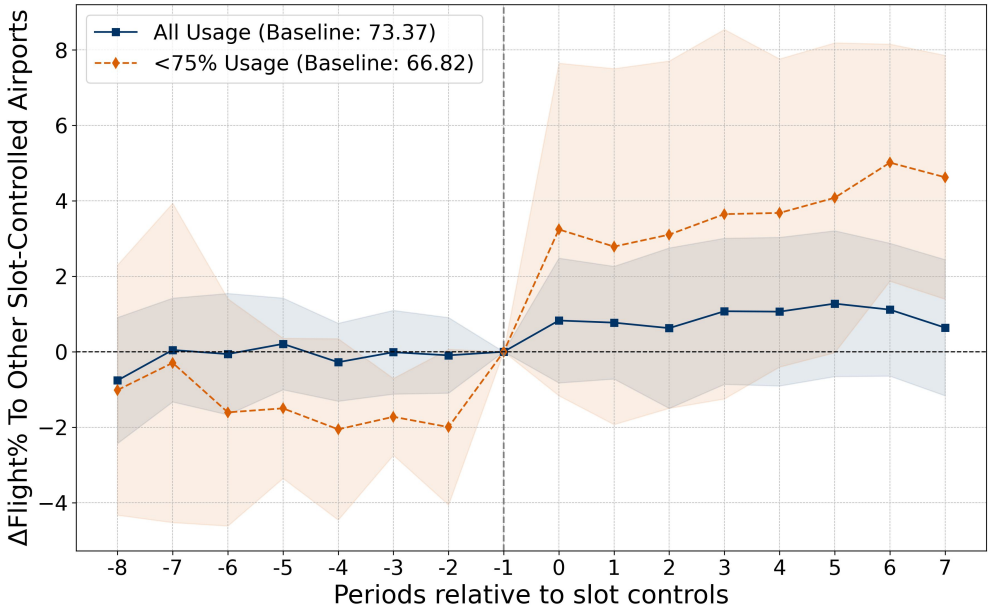


Figure 4: Event study on the percentage of flights destined for other slot-controlled airports. The vertical dashed line indicates the event, and the horizontal dashed red line shows the pre-event baseline average.

An important caveat is that this event study follows eight mid- to small-sized airports immediately after slot controls were imposed. Because these are not Europe's most congested airports, these slots may not be as valuable to airlines. On the other hand, since the study period is right after controls began, most airlines already have a good use for their slots, which were allocated based on the previous year's usage. This explains why overall usage remains constant after slot controls are introduced on these 8 airports. As most airports in Europe have been slot controlled for decades, it is much more likely that in these cases overall usage is higher than it otherwise would be because of the usage requirement.

To better understand the system-wide impacts of slot controls, in the next section I build and estimate a structural model of air travel demand and slot usage in Europe. Given the particularities of this section's 8 treated airports, I do not use the preceding event studies as my main identification strategy for the model.

4 Model

4.1 Framework

The industry has I airlines and C catchment areas. I define an airport's catchment area as all cities within 60km of the airport, in line with Suau-Sanchez et al. (2014). When two airports' catchment areas have cities in common, I combine both airports' catchment areas into a single catchment area.[15]

[15]For the area between the Netherlands and the Rhine-Ruhr region in Germany, applying this definition would result in a very large catchment area with intra-catchment area flights. To avoid this, I divide this catchment area

I define market demand over pairs of catchment areas. Flights on the various routes within catchment areas are thus options the consumer of that market can choose. This aims to model the fact that Europeans frequently travel to neighboring cities to fly, which means that flights from neighboring cities are imperfect substitutes of one another. I exemplify this approach in Figure 5. The picture depicts the Barcelona CA to Paris CA market. Consumers have 3 routes to choose from: they can fly from Barcelona (BCN) to Paris-Charles de Gaulle (CDG), Paris-Orly (ORY) or Paris-Beauvais (BVA). If all routes between catchment areas were perfect substitutes for one another, then the model would not capture the advantages of operating out of congested airports. On the other hand, if each route between catchment areas was an independent market, then the model would not capture the fact that airlines without slots at congested airports can still serve the market by operating from less desirable but less congested airports. My flexible market definition thus allows me to incorporate these trade-offs airlines face when planning routes.

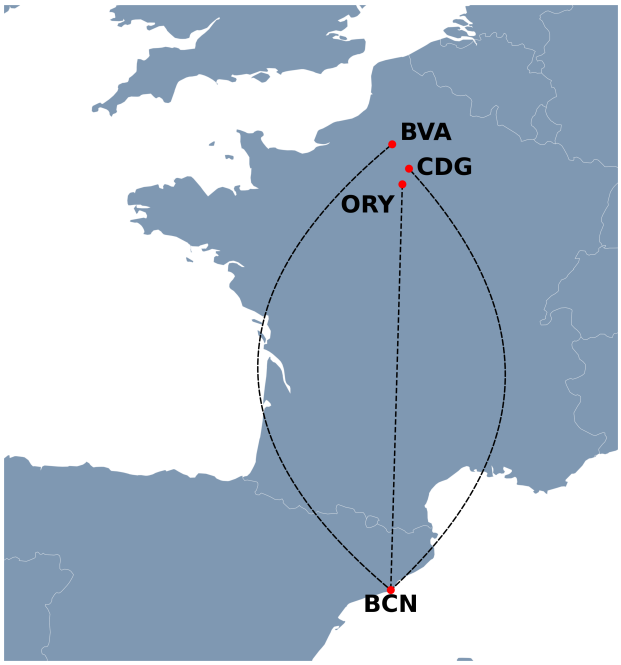


Figure 5: Example of a market between two catchment areas. The picture depicts the Barcelona CA to Paris CA market. Consumers have 3 routes to choose from: they can fly from Barcelona (BCN) to Paris-Charles de Gaulle (CDG), Paris-Orly (ORY) or Paris-Beauvais (BVA). While CDG and ORY have slot controls, BVA does not. If all 3 routes were perfect substitutes of one another, there would be no incentive to own slots at CDG or ORY. On the other hand, if all 3 routes were completely independent, then airlines without slots at CDG or ORY could not compete over consumers in that market. Allowing for imperfect substitution between all 3 routes is thus essential to my model.

Formally, let A_c be the number of airports in catchment area c. Each airport a has a number of slots S_a. Let $S_a = \infty$ denote an airport that has no slot controls. Let S_{iat} denote the number of slots airline i has at airport a during year t. I define a slot allocation $\mathbf{S}_t$ as the vector of S_{iat} for

along national lines.

all airlines i and airports a during period t. I assume that the number of airlines, airports, airport slots, catchment areas and the initial slot allocation $\mathbf{S}_0$ are exogenous. A period in my context is a quarter-year, and I index time by t.

Each directional catchment area pair forms one market. Thus, there are $M \equiv C(C-1)/2$ markets per quarter-year t. Each market may have one or more routes. Let c, c' be the two catchment areas that form market m that consumers choose between. A route in market m is a directional airport pair between one airport of each catchment area. Thus, if catchment area c has $A_c > 1$ airports and catchment area c' has $A_{c'} > 1$ airports, then there are $R_m \equiv A_c \cdot A_{c'}$ routes in market m. Let I_m be the set of airlines that operate in market m, and let $R_{i,m}$ be the set of routes that airline i operates in market m.

Every period (quarter-year) t, airlines first simultaneously choose how often to fly each route. This defines which markets and routes each airline enters. I name flights after their departure hour and day of the week combination, hd. Despite this, daily flights are limited by the number of slots an airline has at both ends of a route.[16] Intuitively, an airline will always need a departure slot at both ends of a route, as the plane has to turn around. Thus, even though I do not explicitly consider arrival slots, they appear in the slot constraints airlines face when choosing where to fly. After determining slot usage, airlines simultaneously compete in prices in all markets they entered.

4.2 Consumer Demand

Each flight in a market represents a different product. Consumers have different valuations for each product based on 5 characteristics: (i) the airline flying the route; (ii) the distance between the two catchment areas; (iii) the airports at each end of the route; (iv) the price; and (v) other product characteristics that the consumer values but that are unobservable to the econometrician. Notice that flights with the same origin and destination catchment area but with different origin and/or destination airports are thus imperfect substitutes of one another.

The indirect utility of a representative traveler that purchases flight hd by airline i for route $r = a \to a'$ in market m during period t is:

$$U(p_{imrt}, \nu^{hd}_{imrt}) = \overbrace{\beta_{p,m} \cdot p_{imrt}}^{\text{price}} + \overbrace{X_{imrt} \cdot \beta_x}^{\text{Other obs charac.}} + \overbrace{\xi_{imrt}}^{\text{Unobs. airline-market charac.}} + \overbrace{\nu^{hd}_{imrt}}^{\sim iid\ T1EV} . \quad (1)$$

Here, p_{imrt} is the average price that airline i charges for flights in route r in market m during period t. I assume that $\beta_{p,m} \equiv \beta_p \cdot \varphi(dist_m)$.[17] Thus, I allow the price coefficient to vary by the distance

[16] In practice, airlines also need to get their slots at origin and destination airports to line up. I abstract away from this problem in this paper for the sake of simplicity. In Appendix D I discuss some strategies that airlines can use to mitigate the slot matching issue.

[17] A more flexible alternative to this would be a full mixed logit model. See Berry et al. (2006), Berry and Jia (2010) and Marra (2024) for applications of such a model to the airline industry. My choice for a simpler model stems from the computational cost of a mixed logit model, as in my model I have 19 airlines making price and frequency choices for as many as 24,000+ routes. This is particularly problematic in Section 8.4, where I employ a genetic algorithm to find the slot allocation that maximizes airline profits as an approximation to the auction outcome. This requires solving the static equilibrium for 64 different slot allocations over 100 generations which takes a long time even with

between the catchment areas in market m. This aims to capture the fact that catchment areas that are closer to each other have better ground transportation, which results in higher price sensitivity.[18]

$X_{imrt} \cdot \beta_x$ is given by:

$$X_{imrt} \cdot \beta_x = \overbrace{\beta_{hub,a'} \cdot hub_{iat} + \beta_{hub,a'} \cdot hub_{ia't}}^{\text{connectivity at orig (a) and dest (a')}} + \overbrace{\eta_{imr}}^{\text{Route·Airline FE}} + \overbrace{\psi_t}^{\text{Quarter·Year FE}} .$$

In addition to airline-by-route and quarter-by-year fixed effects, this vector also includes hub_{iat} and $hub_{ia't}$, which I define as the number of connections that airline i has at both ends of route $r \equiv a \to a'$. These two measurements aim to capture the fact that consumers are more likely to fly with airline i if airline i has more destinations out of airport a and/or into airport a'. This is because they may have already flown with the airline in the past, and/or they may have frequent flier status with the airline.

I assume that the only flight-specific component in the utility function is the T1EV shock ν^{hd}_{imrt}. This assumption means that for each route, airlines only need to choose how many daily flights to fly, instead of needing to choose one flight for each hour-day combination. As in Section 8.4 I need to solve the equilibrium over 6,000 times to approximate the auction outcome, this simplifying assumption allows for reasonable computational times.[19] The fact that I do not explicitly include departure hour/day preferences means that these load onto the route-airline fixed effect. I show evidence of this in Appendix B. Finally, ξ_{imrt} is a collection of product characteristics that the consumer values but that are unobservable to the econometrician.

A consumer chooses between all flights offered by all airlines for a given market and the outside option $i = 0$. In this context, the outside option is taking ground transportation—like trains, buses and ferrys—or not travelling. The consumer's utility from the outside option is:

$$U_{0mrt} = \underbrace{\nu_{0mrt}}_{\sim iid\ T1EV} .$$

The T1EV errors lead to the following expression for the market share of flight hd by airline i on route r in market m during period t with respect to the market share of the outside good:

$$\frac{share^{hd}_{imrt}}{share_{0mt}} = \exp\left(\beta_{p,m\cdot imrt} + X_{imrt} \cdot \beta_x + \xi_{imrt}\right).$$

this relatively simple model.

[18] Alternatively, I could assume that the utility is equal to

$$U(p_{imrt}, \nu^{hd}_{imrt}) = \beta_p \cdot p_{imrt} + \frac{1}{\varphi(dist_m)} \cdot \left(X_{imrt} \cdot \beta_x + \xi_{imrt} + \nu^{hd}_{imrt}\right).$$

Intuitively, this alternative specification states that while all consumers have the same sensitivity to price, the non-price characteristics are more important the longer the flight. This expression is functionally equivalent to (1), as multiplying it by $\varphi(dist_m)$ is a monotone transformation.

[19] Furthermore, In Appendix A I show a restated version of the problem where consumers are allowed to have heterogeneous valuations for flights that depart at different hours and/or days. The problem can be estimated in a similar way to the main text's problem under certain behavioral assumptions on the airline.

Let f_{imrt} be the number of flights airline i offers on route r in market m during period t. Then, the demand equation has the following linear closed form expression:

$$\ln\left(\frac{share_{imrt}}{share_{0mt}\cdot f_{imrt}}\right) = \beta_{p,m}\cdot_{imrt} + X_{imrt}\cdot\beta_x + \xi_{imrt}. \tag{2}$$

4.3 Airline Problem

Each period, airlines make two choices. They first simultaneously choose frequencies, and then they choose prices. Formally, airline i's profits for route r in market m during period t are:

$$\begin{aligned}&\Pi_{imrt}(p_{imrt}, \mathbf{p_{-imt}}, \mathbf{f_{imt}}, \mathbf{f_{-imt}}) = \\ &(p_{imrt} - MC^q_{imrt})\cdot share_{imrt}(p_{imrt}, \mathbf{p_{-imt}}, \mathbf{f_{imt}}, \mathbf{f_{-imt}})\cdot MarketSize_m - MC^f_{imrt}\cdot f_{imrt},\end{aligned} \tag{3}$$

where p_{imrt} is airline i's price in route r; $share_{imrt}$ is airline i's market share from route r $\mathbf{f_{imrt}}$ is airline i's weekly flight frequencies vector for all flights in market m; $\mathbf{p_{-imrt}}$ is a vector with all prices that airline i charges for all other routes in market m, and of all prices that all other airlines charge for all their routes in market m; $\mathbf{f_{-imt}}$ is a vector with all weekly flight frequencies for all routes that airline i offers in market m, and of all flight frequencies that all other airlines offer in all routes of market m; MC^q_{imrt} and MC^f_{imrt} are airline i's marginal costs for adding one extra passenger and one extra weekly frequency, respectively; and $MarketSize_m$ is market m's market size.

Airlines are forward looking and maximize intertemporal profits, taking into account how their current slot usage affects future slot allocations. Specifically, let $V_i(\mathbf{S}_t)$ represent airline i's value function given slot allocation $\mathbf{S}_t$. By the principle of optimality, this value function is implicitly defined by:

$$\begin{aligned}V_i(\mathbf{S}_t) = &\max_{\mathbf{F}_{i,t}}\left\{\Pi_{it}(\mathbf{F}_{i,t},\mathbf{F}_{-it}) + \delta\sum V_i(\mathbf{S}_{t+1})T(\mathbf{S}_{t+1}|\mathbf{S}_t,\mathbf{F}_{i,t},\mathbf{F}_{-i,t})\right\}, \\ &s.t.\sum_{a'\neq a} f^{h,d}_{ia\to a't} \leq S_{iat}, \text{ and} \\ &\sum_{a\neq a'} f^{h,d}_{ia\to a't} \leq S_{ia't} \text{ for all routes } r = a\to a'.\end{aligned} \tag{4}$$

where $\mathbf{F}_{i,t}$ is the vector of all weekly frequencies airline i offers for all routes it participates during period t; $\mathbf{F}_{-i,t}$ is the vector of all weekly frequencies all other airlines $-i$ offer for all routes they participate during period t; $\Pi_{it}(\mathbf{F}_{i,t},\mathbf{F}_{-it})$ are the aggregate stage payoffs airline i gets during period t given frequency choices $(\mathbf{F}_{i,t},\mathbf{F}_{-it})$ (or in other words, the sum of equation (11) for all routes airline i participates in); $T(\mathbf{S}_{t+1}|\mathbf{S}_t,\mathbf{F}_{i,t},\mathbf{F}_{-i,t})$ is the slot transition matrix that determines how slots will change hands during period t given the initial slot allocation during period t $(\mathbf{S}_t)$ and given every airline's frequency choices $(\mathbf{F}_{i,t},\mathbf{F}_{-it})$. I will assume that T is differentiable and that $\frac{\partial T}{\partial f^{hd}_{imrt}} \geq 0$, meaning that increasing frequencies at a route cannot decrease the expected number of slots at that

route's airports in $t+1$.[20] Finally, $\delta \in (0,1)$ is the common discount factor.

Airlines solve problem (4) subject to the slot constraint for each airport in each route $r = a \to a'$. Here, S_{iat} and $S_{ia't}$ are the number of slots that airline i has at airport a and a' during period t, respectively; $\sum_{a' \neq a} f^{h,d}_{ia \to a't}$ is the sum of all of airline i's daily frequencies in all routes that depart from airport a during in d and hour h during period t; and $\sum_{a \neq a'} f^{h,d}_{ia \to a't}$ is the sum of all of airline i's daily frequencies in all routes that arrive at airport a' during in d and hour h during period t. In words, the restrictions require that airline i to not fly beyond its capacity at both the origin and destination airports.

I will assume that the only way an airline can directly affect its future slot allocation $\mathbf{S_{i,t+1}}$ is by choosing to fly a slot less than 80% of the time. I will remain agnostic about all other aspects of the transition matrix dT. In particular, I assume that how airport slots get reallocated is exogenously given, and that airlines know dT. As mentioned in Section 2, in real life the transition matrix will depend on which airlines will request slots to which coordinators, and on how the coordinators prioritize slot requests when dealing with non-grandfathered slots. Exactly pinning down this matrix falls outside the scope of this paper, as this paper's main counterfactual focuses on shutting down the link between today's actions and tomorrow's allocations.

Notice that since airlines choose prices after choosing frequencies, and since their price choices do not affect their slot allocation next period, the airlines' frequency choices fully determine the price game equilibrium. Furthermore, I know that given some set of frequencies, this standard Nash-Bertrand pricing subgame has a unique equilibrium price vector (Mizuno, 2003; Nocke and Schutz, 2018). Thus, I know I can write prices in terms of frequencies.[21]

5 Data

I combine data from three main sources: the list of slot-controlled airports in Europe from IATA; a slot usage dataset based on the flight-level information from Eurocontrol, Europe's air traffic control agency; and the monthly ticket fare and tickets sold data for each direct route in Europe from OAG, an aviation intelligence company.

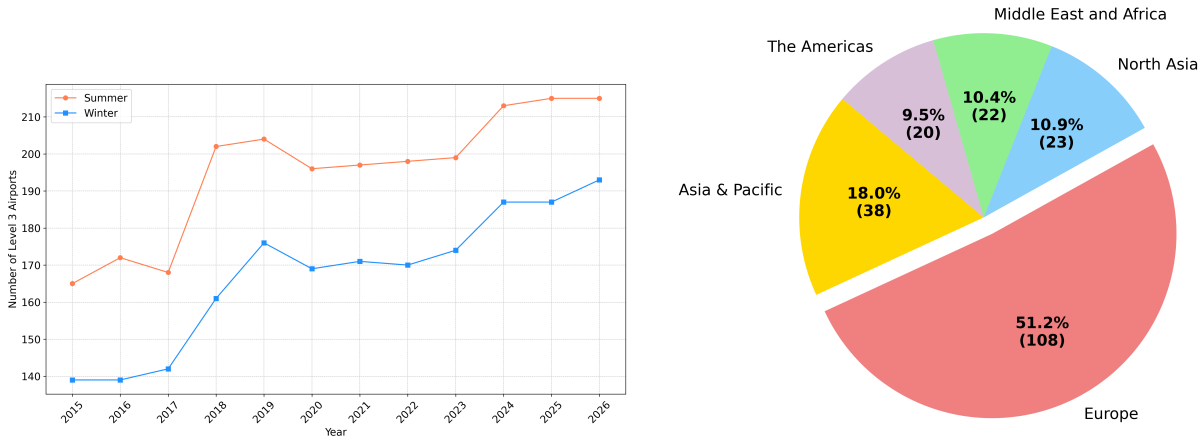
My dataset covers 19 airlines and 156 airports, which together cover over 98% of all intra-Schengen area routes between 2015 and 2019. The decision to include only Schengen flights stems from the fact that passengers in these flights do not need to pass through customs and immigration when traveling. Thus, all airports in Europe divide their available capacity between Schengen and non-Schengen slots, the former of which can only be used to fly to other Schengen flights. This allows me to restrict the airlines' choice sets to a manageable size when performing my counterfactual simulation.

[20] This differentiability rule is not necessary, but I make it for ease of exposition in Section 6.2. Intuitively, this would be equivalent to assuming that there is a probability that each flight gets canceled due to an airline problem.

[21] On the other hand, the entry/frequency choice game may very well have more than one equilibrium. Throughout this paper, the implicit assumption I make is that all simulations and observations come from a unique equilibrium.

5.1 Slot Restrictions by Year (IATA)

I build a dataset of airports with slot restrictions over time based on information from IATA. IATA has on their website the current list of Level 2 and Level 3 airports.[22] I use the Wayback Machine to download the same list for all previous years going back to 2015. In this paper, I focus on Level 3 airports, as these are the airports where an airline needs an arrival/departure slot pair to operate.



(a) Level 3 airports by season, 2015–2026.

(b) Level 3 airports by IATA region, 2019.

Figure 6: Overview of Level 3 (slot coordinated) airports. Some airports are only slot coordinated during either the Summer season (from April to October) or the winter season (from November to March)

Figure 6a shows the evolution of Level 3 airports over time for both the Summer and Winter season, whereas Figure 6b shows the regional share of slot coordinated airports during 2019, the last year I use for my structural model. As I mention in Section 2, all years are divided into two seasons for coordination purposes, a Winter season that runs from November to March, and a Summer season from April to October. Airports usually have the same level for both seasons, but in Europe some airports are Level 3 only for one of the seasons. Figure 6a shows that the number of congested airports has grown over time.[23] On the other hand, Figure 6b shows that the vast majority of congested airports are in Europe. Indeed, 71 of the top 100 busiest airports in Europe are Level 3, including all of the top 25 busiest airports.

5.2 Slot Usage (Eurocontrol's Aviation Data Repository for Research)

I build a dataset of slot usage using Eurocontrol's Aviation Data Repository for Research (ADRR). Eurocontrol is the international organization in charge of coordinating air traffic control (ATC) across the continent. All European countries except Azerbaijan, Belarus and Russia are Eurocontrol

[22] The current list can be found in Annex 12.7 of the Worldwide Airport Slot Guidelines (IATA, 2025).

[23] The big spike in 2017-2019 mainly responds to 16 Brazilian airports and 10 Chinese airports that started having slot controls during this time period. Both the Brazilian and Chinese civil aviation authorities changed the rules on how to measure airport capacity around this time period, in order to align themselves with the WASG. For more information, see Resolution 440/2017 of the *Âgencia Nacional de Aviação Civil* (ANAC) for the Brazilian case, and Fu et al. (2020) for the Chinese case.

members. The ADRR is a dataset that has flight-level data for all flights that crossed the airspace of Eurocontrol member states. For each flight, the dataset includes the origin and destination airports, scheduled and actual departure and arrival times, flight operator, aircraft type, registration number, distance flown, flight path and market segment (Lowcost, Mainline, Cargo, Charter, Business). This data is available for the months of March, June, September and December, between 2015 and 2021, and it includes information on over 19.6 million flights. I use this dataset to reconstruct the slot series. The details of this process can be found in Appendix E.

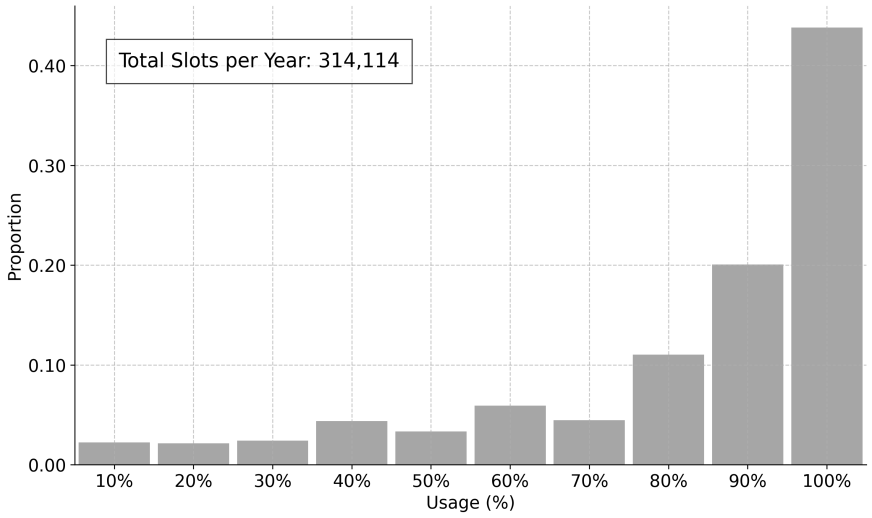


Figure 7: Histogram of slot usage. I name bins based on the center value of the bin. Thus, the 80% bin includes all slots used between 75% and 85%.

Figure 7 shows the Schengen area slots divided into 10% bins.[24] As can be seen, the vast majority of slots are used intensively. By contrast, only 10% of slots are used around the 80% threshold. As most slot-constrained airports have had capacity issues for many years, most airlines should have a good use for their slots. That being said, there could still be allocative inefficiencies: while airline i currently has a good use for slot s, airline i' may have an even better use for the slot, in terms of either total profits and/or total consumer surplus. The total number of slots over the years has gone up, both because of marginal increases in capacity, as well as because of the 8 airports that become slot constrained over this period.

5.3 Route Prices and Passengers (OAG's Travel Analyzer)

To estimate the air travel demand for the Schengen Area, I augment the ADRR data with prices and quantities data from OAG's Travel Analyzer. OAG is an aviation intelligence company, and their dataset includes monthly average fares and total passengers for each route-airline combination in the Schengen area. I aggregate these data up to the quarter level, as the ADRR data has only one month per quarter available.

[24] See Appendix I for a detailed breakdown of slot ownership by airline.

Group	Passengers (Mn)	Revenue (Bn $)	Avg. Fare ($)	Monopoly Share
Lufthansa Group	397.9	58.01	145.79	44.0%
Air France-KLM	294.2	41.09	139.69	33.4%
Ryanair	596.6	31.58	52.94	40.8%
EasyJet	345.4	25.39	73.51	21.4%
IAG (BA-Iberia)	290.6	25.35	87.25	27.4%
Norwegian	187.1	19.94	106.56	32.4%
SAS Group	142.1	16.48	115.98	21.8%
Air Berlin	80.3	11.77	146.49	27.2%
Alitalia	87.1	8.93	102.48	13.8%
TAP Air Portugal	56.1	5.41	96.46	20.4%
Aegean Airlines	56.7	4.14	72.94	42.1%
Finnair	36.1	3.89	107.73	53.1%
Wizz Air	72.7	3.68	50.67	45.1%
Air Europa	31.1	3.62	116.59	27.9%
Volotea	34.7	2.81	80.89	54.5%
LOT Polish Airlines	20.3	2.54	125.28	32.8%
Luxair	8.3	1.43	171.59	37.6%
Air Malta	9.7	1.14	117.85	25.7%
Czech Airlines	12.4	0.99	80.23	23.1%

Table 1: Summary of key performance indicators by airline group, 2015-2019. The dataset includes the top 19 airlines by revenue in the Schengen Area, who collectively encompass over 98% of all Schengen Area revenue.

OAG gets their data from Travelport, one of the 3 main global distribution systems (GDS) companies. Along with Amadeus and Sabre, Travelport provide airlines, hotels and travel agencies with the computerized network systems to process bookings. Thus, all 3 companies have access to airline ticketing data. OAG uses Travelport data as their input to construct estimates for monthly average fares and tickets-sold data using proprietary algorithms. The fare data that Travelport supplies does not include tax data, but as airfare taxes are relatively uniform across the Schengen area, this is not a major concern.

Table 1 presents some summary statistics for all the airline groups in my data. From the data we can see that the top 5 airline groups (AirFrance-KLM, EasyJet, IAG, Lufthansa and Ryanair) represent around 67% of total Schengen area revenue and 69% of total Schengen area passengers. We can also see variation in the average fare airlines charge, and also in the share of monopolist routes.

I use the OAG dataset to define the bases for each airline group. An airport is a base for airline i if i flies at least 10 routes out of that airport. Table 2 presents a summary of all groups' bases.[25] We can see that my definition of base covers the vast majority of routes for all airlines. In addition, we can also see that largest airline groups have on average more bases than small airlines, and that Low-cost groups have more bases than full-service carriers. This last point falls in line with Low-cost carriers' point-to-point model, compared to full-service carriers' hub-and-spoke model. In my simulations, all routes that an airline offers have to either begin or end at one of its bases.

[25] See Appendix I for a detailed breakdown of all bases for each airline group.

Airline Group	Total Bases	Hub Pax %	Top Hub 1	Top Hub 2
Ryanair	79	100.00%	Milan (BGY)	Barcelona (BCN)
Lufthansa Group	30	99.98%	Munich (MUC)	Frankfurt (FRA)
EasyJet	26	100.00%	Milan (MXP)	Geneva (GVA)
IAG (BA-Iberia)	26	99.85%	Barcelona (BCN)	Madrid (MAD)
Volotea	18	95.12%	Venice (VCE)	Nantes (NTE)
Air France-KLM	16	99.60%	Amsterdam (AMS)	Paris (ORY)
Norwegian	15	99.84%	Oslo (OSL)	Stockholm (ARN)
Air Berlin	13	99.36%	Berlin (TXL)	Düsseldorf (DUS)
Wizz Air	7	90.94%	Budapest (BUD)	Gdańsk (GDN)
Air Europa	6	95.59%	Madrid (MAD)	Mallorca (PMI)
SAS Group	5	88.38%	Oslo (OSL)	Stockholm (ARN)
Aegean Airlines	4	99.74%	Athens (ATH)	Thessaloniki (SKG)
Alitalia	3	95.28%	Rome (FCO)	Milan (LIN)
TAP Air Portugal	2	99.97%	Lisbon (LIS)	Porto (OPO)
LOT Polish Airlines	1	88.67%	Warsaw (WAW)	N/A
Czech Airlines	1	97.73%	Prague (PRG)	N/A
Luxair	1	100.00%	Luxembourg (LUX)	N/A
Air Malta	1	99.19%	Malta (MLA)	N/A
Finnair	1	99.46%	Helsinki (HEL)	N/A

Table 2: Base summary statistics by airline group, 2015-2019. Airline Bases include the IATA airport code.

5.4 Other Data Sources

As I explain in Section 4, I define a market as a catchment area (CA) pair. An airport's CA includes all people living in cities within 60 km from the airport. When two cities are part of two or more airport's CAs, then all of the cities in both airport's CAs form part of the same CA. In Appendix I I detail all CAs with multiple airports. I take the population data from Eurostat, and I include any urban agglomeration with more than 30,000 inhabitants in 2017.

Finally, for my instrumental variables I use the jet fuel price by year, which I take from the US Energy Information Administration's U.S. Gulf Coast Kerosene-Type Jet Fuel Spot Price FOB series; the overflight rights unit price, which I take from Eurocontrol's Annual evolution of en-route Unit Rates dataset; the exchange rate of European currencies and the US Dollar, which I take from the World Bank's Global Economic Monitor; and the aircraft maximum take off weight (MTOW), which I take from Eurocontrol's Aircraft Performance Database.

6 Estimation Strategy

6.1 Demand

Equation (2) is a linear equation with prices as the sole endogenous variable. Prices are endogenous because of the classic endogeneity associated with the unobserved demand shock ξ_{imrt}. I employ two cost-shift instruments for estimation: fuel cost per passenger, and overflight dues per passenger. The former is due to Marra (2024),[26] whereas the latter is similar in spirit but to the best of my

[26]Previously, the literature also used fuel-by-distance as a proxy of this instrument (Berry and Jia, 2010).

knowledge it is novel to my paper. In both cases, I rely on my dataset's panel structure to isolate the residual variation in fuel prices and overflight dues after controlling for airline-route and period fixed effects. These two instruments have the advantage that they are route- and airline-specific, as they both depend on the distance flown as well as the aircraft type the airline uses for a given route.

My first instrument is the fuel cost per passenger. As the Eurocontrol R&D Dataset does not include fuel consumption, I estimate each individual flight's fuel consumption using the Fuel Estimation in Air Transportation (FEAT) model that Seymour et al. (2020) develop. The FEAT model computes fuel consumption as a quadratic formula of distance, with aircraft-specific parameters. Specifically, a flight with aircraft type i will consume an amount of fuel F_i according to the following formula:

$$F_i = \alpha_i \cdot d_{\text{gc}}^2 + \beta_i \cdot d_{\text{gc}} + \gamma_i,$$

where d_{gc} is the great-circle distance between origin and destination, and α_i,β_i and γ_i are aircraft-specific parameters that the authors estimate. Seymour et al. (2020) show that the model has an error of less than 5% when compared to actual fuel burn reports.[27]

My second instrument is the overflight dues per passenger. Whenever an aircraft overflies a country, its owner has to pay that country overflight dues to help cover ATC costs. In Europe, these dues are governed by the Eurocontrol treaty.[28] The key provision in this treaty is that the overflight dues that an airline owes to a country are based on the aircraft's size, the distance covered over the country, and a country specific multiplier. Formally, the overflight dues that flight f by airline i flying from a to a' pays during period t are:

$$\text{Overflight dues}_{fia\to a'mt} = \sum_{c\in C} \text{unit price}_{ct} \times \frac{\text{distance (km)}_{caa'}}{100} \times \sqrt{\frac{\text{Max Take-Off Weight}_{fia\to a'mt}}{50}}.$$

Here, Max Take-Off Weight is the maximum weight an aircraft can have at take-off. I take this data from Eurocontrol's Aircraft Performance Data. On the other hand, unit price is the country-specific multiplier, which I also take from Eurocontrol. Finally, I compute the distance based on the countries covered by the great-circle distance between a and a'.[29] The total overflight dues are a sum of the overflight dues that each overflown country charges. Eurocontrol sets the unit price so that the national ATC provider breaks even. The price is set once per year in November, and it is set in the country's national currency. This fact coupled with currency fluctuations gives me the variation needed to instrument airfare prices.

[27] This error comes from various flight specific characteristics, like altitude, aircraft speed, aerodynamic drag, to name a few. There exist more sophisticated models that take these into account, like Eurocontrol's Base of Aircraft Data (BADA) model (Nuic et al., 2010). While this extra level of accuracy is important for many aeronautical and engineering applications, the FEAT model is accurate enough for my purposes.

[28] Further details on the treaty provisions pertaining to overflight due payments can be found at https://www.eurocontrol.int/publication/customer-guide-route-charges. Accessed on July 10, 2025.

[29] In practice, airlines frequently deviate from this shortest path to avoid expensive airspace or delay-prone sectors. I abstract away from this fact and use the shortest path between a and a' for both instruments.

In both cases, the identifying assumption is that the residual variation in overflight dues and fuel prices after controlling for route-airline (id) and quarter-year (time) fixed effects is uncorrelated with the demand shocks. This requires for changes in country unit rates not to be correlated with route-specific demand shocks for routes that overfly that country. While I cannot completely rule out this concern, the fact that unit rates are chosen based on break-even objectives for the ATC agency adds substantial variation that should be uncorrelated to demand shocks.[30] On the other hand, we might be worried about the fact that both cost-shifters are aircraft specific. This will not be a problem if two things are true: first, we need to assume that consumers do not observe or care about their aircraft's size or fuel efficiency, an assumption that seems warranted for the 2015-2019 period that I study. Second, it needs to be true that airlines choose the aircraft type before the demand shock is revealed. This seems warranted given that in Europe airlines need to inform coordinators about their flight's specifications months in advance, as explained in Section 2.

For estimation, I assume that

$$\varphi(dist_m) = 1 + \gamma_1 \cdot dist_m$$

This function will be decreasing in $dist_m$ depending on the value of γ_1. The demand equation from Equation (2) will thus be:

$$\ln\left(\frac{share_{imrt}}{share_{0mt} \cdot f_{imrt}}\right) = \beta_p \left(1 + \gamma_1 \cdot dist_m\right) \cdot_{imrt} + X_{imrt} \cdot \beta_x + \xi_{imrt}. \tag{5}$$

An additional assumption to correctly identify (5) is that airlines choose frequencies f_{imrt} before the demand shocks ξ_{imrt} are realized. Intuitively, European airlines have to decide their schedules months in advance when they inform coordinators of their flight specifications, and this time-to-build assumption falls in line with the literature on frequency choice for airlines (Aguirregabiria and Ho, 2010; Yuan and Jia-Barwick, 2024).

To avoid potential measurement errors in my data affecting my results through extreme values, I restrict my demand estimation sample to those markets where there were at least 260 passengers in the quarter, were the average fare was at least \$15.[31] In addition, I also remove markets where the OAG prices are too low given the estimated fuel and overflight dues per passenger. I do this to address potential measurement error in OAG's price data, as my estimates for the fuel and overflight rights costs would make flying that route deeply unprofitable at those prices. These removed routes make up around 6.15% of the total sample size.[32]

Finally, I assume that the market size between CA c and CA c' is equal to the geometric mean of both CAs' *adjusted* population.[33] I adjust the population in two ways. First, I multiply each CA's

[30] In Figures 21 and 22 of Appendix I I show the evolution of fuel prices and unit rates over 2015–2019.

[31] These restrictions are similar to those frequently imposed on the US Department of Transportation's DB1B database. See, e.g., Berry and Jia (2010); Aguirregabiria and Ho (2010).

[32] Specifically, I remove routes where fuel costs and overflight costs per passenger represent more than 44.58% and 17.65% of the ticket price, respectively.

[33] This approach is similar to the literature's approach, where market size usually equals the geometric mean of the two ends' populations (Berry et al., 2006).

population by the log difference between the country's 2017 income and the average EU income, which I also take from Eurostat. This aims to capture the fact that wealthier people travel more. Second, I define all Mediterranean and island CAs with less than 500,000 people as "tourism-heavy" CAs. When a market includes one "tourism-heavy" CA and one regular CA, I assume that the market size is equal to the regular CA's population. This aims to capture the fact that the vast majority of passengers in these pairings are tourists hailing from the regular CA.

6.2 Marginal Costs

I estimate both the marginal cost of selling one additional ticket (MC^q_{imrt}) as well as the marginal cost of adding one additional weekly flight (MC^f_{imrt}) from the first-order conditions of the pricing and frequency choice problems, respectively. I further decompose these marginal costs on a set of variables and fixed effects. I do this both to account for the effect of intercontinental flights on the scheduling of full-service carriers, and to construct marginal costs for my counterfactual simulations.

For the marginal cost of serving passengers, I use my demand estimates under the Nash-Bertrand assumption to recover the marginal costs. Specifically, from Equation (11) we have that the first order condition is given by:

$$MC^q_{imrt} = \frac{share_{imrt}}{\frac{\partial share_{imrt}}{\partial p_{imrt}}} + p_{imrt}. \tag{6}$$

Given the T1EV assumption, the right-hand side of equation (6) can be directly imputed from the data and the estimated values for equation (5).

For the marginal cost of adding weekly frequencies, I also use my demand estimates. However, since this period's frequencies affect next period's slot allocation, I cannot estimate this marginal cost for all markets. To see why, notice that from the dynamic problem in Equation (4) and the stage profits in equation (11) we have that if airline i was completely myopic about the 80% usage rule and the slot restrictions, it would choose frequencies such that:

$$\sum_{\tilde{r} \in R_m} \left(p_{im\tilde{r}t} - MC^q_{im\tilde{r}t}\right) \cdot \frac{\partial share_{im\tilde{r}t}}{\partial f^{h,d}_{imrt}} \cdot MarketSize_m = MC^f_{imrt}. \tag{7}$$

Here, the left hand side is the marginal revenue of adding one extra flight in route r. Notice that this extra flight will affect the profits of not only route r, but also all other routes $\tilde{r}$ that airline i offers in market m.[34]

[34] As I am backing these out from the FOC in equilibrium, we can use an envelope theorem type argument to ignore the indirect effects of f^{hd}_{imrt} on $p_{im\tilde{r}t}$.

However, if an airline uses a slot 80% of the time, then we have that:

$$\sum_{\tilde{r}\in R_m}\left(p_{im\tilde{r}t}-MC^q_{im\tilde{r}t}\right)\cdot\frac{\partial share_{im\tilde{r}t}}{\partial f^{h,d}_{imrt}}\cdot MarketSize_m+\delta\sum V_i(\mathbf{S}_{t+1})\frac{\partial T(\mathbf{S}_{t+1}|\mathbf{S}_t,\mathbf{F}_{i,t},\mathbf{F}_{-i,t})}{\partial f^{hd}_{imrt}}=MC^f_{imrt}$$

$$\sum_{\tilde{r}\in R_m}\left(p_{im\tilde{r}t}-MC^q_{im\tilde{r}t}\right)\cdot\frac{\partial share_{im\tilde{r}t}}{\partial f^{h,d}_{imrt}}\cdot MarketSize_m\leq MC^f_{imrt},$$

where I am using the fact that $\frac{\partial T(\mathbf{S}_{t+1}|\mathbf{S}_t,\mathbf{F}_{i,t},\mathbf{F}_{-i,t})}{\partial f^{hd}_{imrt}}\geq 0$.

Intuitively, this means that airline i could increase its stage profits by flying one less flight in route r. Doing so could however mean that airline i loses the slot in the future. Thus, by flying exactly 80% of the time, airline i is taking into account the dynamic consequences of today's frequency choice, as stated in Problem (4). Estimating the marginal cost through (7) risks underestimating the true marginal cost.

On the other hand, if an airline i uses all slots in route r exactly 100% of the time, the total slot restriction on that airport might be operative. Specifically, and letting $\lambda_a, \lambda_{a'}$ be the two multipliers for route $r\equiv a\to a'$, we may have that:

$$\sum_{\tilde{r}\in R_m}\left(p_{im\tilde{r}t}-MC^q_{im\tilde{r}t}\right)\cdot\frac{\partial share_{im\tilde{r}t}}{\partial f^{h,d}_{imrt}}\cdot MarketSize_m-\lambda_a-\lambda_{a'}=MC_{imrt}$$

$$\sum_{\tilde{r}\in R_m}\left(p_{im\tilde{r}t}-MC^q_{im\tilde{r}t}\right)\cdot\frac{\partial share_{im\tilde{r}t}}{\partial f^{h,d}_{imrt}}\cdot MarketSize_m\geq MC_{imrt},$$

where I am using the fact that $\lambda_a, \lambda_{a'}\geq 0$.

Intuitively, this means that adding one additional flight would be profitable for airline i, but the slot constraint prevents it from doing so. Estimating the marginal cost through (7) risks overestimating the true marginal cost. To avoid these problems, I construct MC^f_{imrt} using (7) for routes where no slot is used around 80% of the time, and where at least one slot is not used 100% of the time. Then, I use these estimates to project MC^f_{imrt} for the routes where (7) does not hold.

I decompose both marginal costs as:

$$y_{imrt}=\overbrace{\mu^{o,US}_{imrt}\cdot own^{US}_{imrt}+\mu^{o,RoW}_{imrt}\cdot own^{RoW}_{imrt}}^{i\text{'s own intercontinental connectivity}}+\overbrace{\mu^{p,US}_{imrt}\cdot partner^{US}_{imrt}+\mu^{p,RoW}_{imrt}\cdot partner^{RoW}_{imrt}}^{i\text{'s partner intercontinental connectivity}}+$$
$$\underbrace{\mu_{hub,a'}\cdot hub_{iat}+\mu_{hub,a'}\cdot hub_{ia't}}_{\substack{\text{european connectivity at}\\ \text{orig (a) and dest (a')}}}+\underbrace{\mu_{d1}\cdot d_r+\mu_{d2}\cdot d_r^2}_{distance}+\underbrace{\eta_{ia}+\zeta_{ia'}}_{\substack{\text{Airline}\times\\ \text{Airport FE}}}+\underbrace{\psi_t}_{\substack{\text{Airline}\times\\ \text{Quarter}\times\\ \text{Year FE}}}. \quad (8)$$

Here y_{imrt} equals either MC^q_{imrt} or MC^f_{imrt}. Besides variables used in the demand estimation, I also use measures of intercontinental connectivity. These measures correspond to the number of compatible intercontinental flights that airline i has at both origin and destination airports for flights that cover route r. I define compatible intercontinental flights as intercontinental flights that either depart within 1 and 5 hours of the arrival of a flight in route r, or that arrive between 1 and 5 hours

before the departure of a flight in route r.[35]

These measures aim to capture how a full-service carrier's intercontinental network shapes their intra-Schengen scheduling and pricing decisions. For the marginal cost of serving passengers, the measures aim to capture the opportunity cost that full-service carriers have when selling a ticket for Intra-Schengen passengers. This opportunity cost arises because these airlines could have sold that ticket as part of a (more lucrative) intercontinental flight with one connection. For the marginal cost of adding frequencies, the intercontinental compatibility measure aims to capture the fact that intercontinental flights work as an implicit subsidy for full-service carriers when they choose to add an extra intra-Schengen service. This is because part of the costs will be recouped as additional connecting passengers to/from intercontinental destinations.

As in my sample around 57% of routes to or from slot controlled airports have 100% usage for all slots, my imputation for markets where (7) does not hold raises a selection bias. Indeed, these 100% usage routes skew to larger, more congested airports, which means that omitting them leads to a underestimation of the true marginal costs for these routes. However, as this bias leads airlines to fly more to these very congested airports regardless of the usage requirements, it actually makes my counterfactual simulations more similar to the baseline scenario with the 80% requirement. Thus, this bias makes the baseline and counterfactual more similar, which means that my results are actually a lower bar to the true effects of removing the usage requirement.

Still, I partially address this issue by decomposing MC^f_{imrt} using the individual slot series as my units of observation instead of the routes. This is equivalent to applying weights to the route-level decomposition based on the number of slots used for the route.[36] As the 100% usage routes bias towards busier routes, this weighted decomposition should partly address the potential selection bias in my sample.

7 Empirical Results

7.1 Demand

Table 3 presents the results for the demand parameters from equation 5 applying the instruments from Section 6.1. The Kleibergen-Paap (Kleibergen and Paap, 2006) F-stat of 33.55 and the CLR test (Andrews et al., 2019) of 107.44 indicate strong instruments.[37] Given these parameters, we can see that γ_1 is negative, which means that φ is decreasing. The median demand elasticity is 4.43.

[35] I divide these measures into 4 groups. compatible flights that airline i flies to/from the US (i) and to/from the rest of the world (ii), and compatible flights that airline i's non-European alliance partners fly to/form the US (iii) and to/from the rest of the world (iv). A full list of non-European partners for each European full-service carrier can be found in Appendix I. The distinction between the US and the rest of the world stems from the fact that the US market is the most important intercontinental market for European full-service carriers. For example, in 2024 the US accounted for 56% of Lufthansa's non-European revenue (Lufthansa Group, 2024), and 47% of IAG's (IAG, 2024). On the other hand, I include partner flights to account for the code-sharing agreements that most European full-service carriers have with their non-European allies.

[36] To see why, notice that if airline i is using two slots to fly from Airport a to Airport a', then I would have two first order conditions for the Airport a to Airport a' market. then it would be as if airline i's route from a to a' has twice the weight in the decomposition regression for MC^f_{imrt}.

[37] See Table 17 in Appendix I to see the first stage regressions.

This is slightly higher than what previous studies have found for Europe (IATA, 2007; Marra, 2024), and it probably responds to the fact that these previous studies had both intra-European as well as extra-European flights, the latter of which tend to be less sensitive to price as there are fewer alternatives to flying. On the other hand, I find that consumers value if an airline connects them to more destinations at both the origin and destination airport. A 50% increase in the number of connections at an airport leads to an increase of between $5.80 and $19.67 in the average consumer's willingness to pay.

	Ln(Demand)
Fare, in \$100 (β_p)	-6.844***
	(1.614)
Fare, in \$100 · Distance, in 100 km ($\gamma_1 \cdot \beta_p$)	0.159***
	(0.0467)
# Connections at Origin	0.00552
	(0.00337)
# of Connections at Destination	0.0187***
	(0.00617)
Observations	79,356
Kleibergen-Paap F-stat	33.55
CLR test	107.44

*** p<0.01, ** p<0.05, * p<0.1

Table 3: Results from Estimating Demand Equation (2). The regression includes Fixed effects at the Origin×Dest×Airline level and Quarter×Year level. Errors are clustered at the Market (Origin Catchment Area×Dest Catchment Area) level.

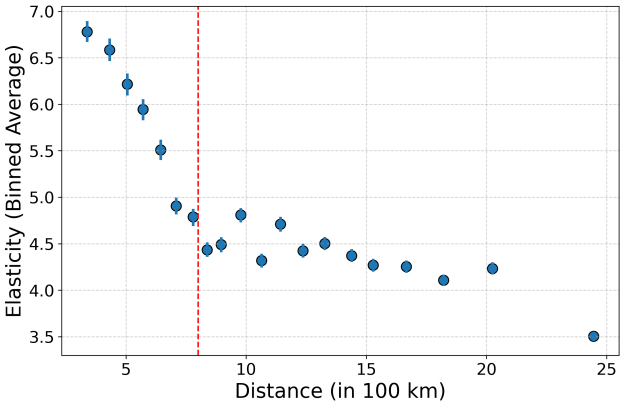


Figure 8: Binscatter of elasticity in terms of distance. I divide routes into 20 bins. Each bin includes the 95% CI for the mean, which I compute through bootstrapping.

Figure 8 shows a binscatter plot of elasticity in terms of distance. I find that farther away markets tend to be less elastic. In general, distance translates into worse ground travel alternatives. For example, the most elastic markets for Frankfurt are Zurich and Munich, both of which can be

reached by train in under 4 hours. By contrast, the least elastic markets for Frankfurt are Lanzarote and Las Palmas, both islands with no ground connections to Frankfurt. In addition, we see that the relationship becomes weaker after 750 km. This falls in line with the common notion that ground transportation is a competitive alternative in Europe for distances below 800 km (IATA, 2003).

7.2 Marginal Costs

Table 4 presents the results of the decomposition estimation for the marginal cost of serving passengers (MC^q_{imrt}, Column 1) and the marginal cost of adding weekly frequencies (MC^f_{imrt}, Column 2). As I explain in Section 6.2 I estimate MC^q_{imrt} from the first order condition of the airline's pricing problem. On the other hand, I estimate MC^f_{imrt} by restricting attention to routes where the airline has no slot at 80% usage, and where at least one slot is at less than 100% usage. As the cost distribution for MC^f_{imrt} has a long right tail, I take logs for the decomposition.

The average marginal cost of serving passengers is \$78. Low-cost carriers have a marginal cost of \$52, while full-service carriers have a marginal cost of \$98. I find that this marginal cost is convex in distance, although the quadratic coefficient is not significant. On the other hand, I find that this marginal cost is decreasing in the number of connections at the origin airport. This may reflect the economies of scale associated with airline bases, like more efficient ticketing and better airport agreements on per passenger airport charges.

I find that the intercontinental connections to and from the US increase the opportunity cost of selling a ticket for intra-Schengen flights. Predictably, partner connections lead to a lower opportunity cost increase than own connections. As the average full-service carrier route in my data has 1.74 daily own US connections and 1.54 daily partner US connections, the total effect of this opportunity cost is generally very small. However, it is significant for several of the busiest full-service carrier routes. For example, the opportunity cost due to US connections represents 37% of the marginal cost for AirFrance-KLM on the Amsterdam to Paris route, and 38% for Lufthansa on the Frankfurt to Munich route. The connections from the rest of the world do not appear to be as critical.

On the other hand, the average marginal cost of adding one additional weekly frequency is \$394k per quarter. This cost is \$520k for low-cost carriers and \$355k for full-service carriers. I find that this cost is increasing in distance and also in the number of connections at the origin airport. This may reflect the fact that some airport-specific operations become increasingly complicated with scale.

Intercontinental connections decrease the perceived cost for flying intra-Schengen flights. The effect is largest for own intercontinental connections to the US and to a lesser extent to the rest of the world. Partner connections on the other hand do not appear as critical. Again, the low average daily connections to the US and the rest of the world for most of my sample mean that the total effect of this implicit subsidy is generally very small.[38] However, it is again significant for several of the busiest full-service carrier routes. For example, the implicit subsidy for full-service carriers when they choose to add an extra intra-Schengen service is equal to 82% of the actual cost

[38] To compute this implicit subsidy, I use my estimates to predict the marginal cost for a given route if all intercontinental connections were set to 0.

	(1) MC^q_{imrt}	(2) $\ln(MC^f_{imrt})$
Distance (in 100 km)	2.013***	0.0751***
	(0.246)	(0.0159)
Distance (in 100 km)^2	0.0111	-0.00222***
	(0.00755)	(0.000535)
# Connections at Origin	-0.0823*	0.00686*
	(0.0426)	(0.00390)
# Connections at Destination	0.0573	0.00208
	(0.0431)	(0.00387)
# Own Intercontinental Connections, US	0.548**	-0.0146**
	(0.277)	(0.00677)
# Own Intercontinental Connections, non-US	-0.0709	-0.0102***
	(0.0660)	(0.00296)
# Partner Intercontinental Connections, US	0.227	-0.00173
	(0.154)	(0.00518)
# Partner Intercontinental Connections, Non-US	-0.0413	-0.00698
	(0.145)	(0.00675)
Observations	111,072	242,858
R-squared	0.673	0.491
Unit	Route	Slot

*** p<0.01, ** p<0.05, * p<0.1

Table 4: Results from marginal costs decomposition. The regression includes Fixed effects at the Origin×Airline, Dest×Airline and Quarter×Year×Airline level. Errors are clustered at the Market (Origin Catchment Area×Dest Catchment Area) level. The Flight MC is in logs because of the long right tail in the Flight MC distribution.

for AirFrance-KLM on the Amsterdam to Paris route, and 48% for Lufthansa on the Frankfurt to Munich route.

8 Counterfactual Simulations

8.1 Baseline Simulation Strategy

The dynamic model that I set up in Section 4 is a dynamic version of a Generalized Assignment Problem (GAP), where the agents are the airports, the tasks are the flights to be scheduled, and the constraints are the number of slots at each airport. As the GAP is an NP-hard problem (Özbakir et al., 2010), solving for my model's equilibrium without imposing additional behavioral assumptions is computationally challenging. As I need a solution under grandfathering as a baseline to compare my counterfactual simulations to, in this section I present a simplifying assumption on airline behavior that allows me to approximate for the airline's equilibrium choices under grandfathering.

To bypass my problem's computational challenges, I impose the behavioral assumption that

under grandfathering airlines behave as though as they are required to use their slots at least 80% of the time, and they believe all other airlines behave the same way. Operationally, this turns the dynamic problem in Section 4 into a static problem with a set of usage restrictions for slot-constrained airports. This assumption is a stylized simplification of the actual decision process airlines face when choosing which routes to fly. As airlines have been known to fly empty flights to keep their slots,[39] and as most slots at the busiest airports are used at least 80% of the time in my data, this stylized baseline should be a reasonable approximation to actual airline behavior. In Section 8.2 I show that the computed baseline aligns well with the share of monopoly routes event study in Section 3.

Assumption 1. *Assume airport a is slot-controlled. Then, for all $i \in I$ we have that airline i behaves as though as they face the following restriction:*

$$\sum_{a' \neq a} f^{h,d}_{ia \to a' t} \geq S_{iat}$$

where the left hand side is the sum of all flights airline i has at airport a, and the right hand side is the total number of slots airline i has in airport a. Furthermore airline i believes all other airlines face the same restriction.

Still, we might expect some slots to be more valuable than others. To account for this, in Appendix F I recompute the baseline when airlines behaved as though as they have to keep slots at only the busiest airports, or if they only have to keep certain slots. The results do not meaningfully change with these alternative assumptions. I favor the 80% over all slots for the main text because it is a simpler heuristic than the alternative specification I propose in Appendix F.

On the other hand, the fact that as I mention in Section 4 the hour/day preferences load onto the route-airline fixed effects means that the demand has a potential bias. To see why, notice that if i's flight from a to a' had a particularly good departure time, then my route-airline fixed effects will indicate that this route-airline combination is very desirable *for all of i's flights from a to a'*. The net effect of this bias is unclear, as on the one hand incentivizes airlines to concentrate their flights on fewer routes, but it also makes those routes artificially desirable to consumers. To address this concern, in Appendix B I perform a correction to the route-airline FE where I replace the hour/day-specific effect with an average effect. This change leads to even stronger results, mainly due to a bigger drop in consumer welfare.

8.2 Baseline: Grandfathering with Current Slot Allocation

I first construct a baseline simulation that serves as a comparison to my counterfactual simulations. As I mention in Section 8.1, the dynamic problem in Section 4 is NP-hard, which means that I need to impose additional assumptions to simulate it. I avoid the curse of dimensionality through a stylized approximation of airline behavior, where I assume they behave as if they are required to

[39] For example, Lufthansa's CEO Carsten Spohr stated in 2022 that: "We have to operate 18,000 additional, unnecessary flights during the winter, purely to secure our slots." (Kösters et al., 2023)

keep all their slots. The equilibrium from this stylized behavior matches the size for the monopoly routes event study in Section 3.

To simulate the equilibrium, I first construct cost estimates for all markets that no carrier currently serves. From the event study in Section 3 we know that airlines change their networks in response to changes in the slot allocation rules, so having estimates for these markets allows me to better simulate the counterfactual network changes. Notice that European connectivity is an endogenous object, which will result from the airlines' route choices. On the other hand, I compute the population-weighted great-circle distance between all catchment areas, which serves as input to both the costs and the demand. I take the intercontinental connectivity as exogenous and airport-period specific. This intercontinental connectivity only affects existing routes, which aims to capture the fact that the existing routes already have a robust intercontinental demand. Finally, I assume that whenever an airport-airline pair does not have a cost fixed effect, then that airport-airline pair's fixed effect is equal to the highest fixed effect among all airlines that operate at that airport. This aims to capture the fact that if an airline doesn't operate at an airport, it must be particularly costly for that airline to operate there.

Next, I construct demand estimates for all non-existent markets. I estimate the distance and European connectivity measures in a similar fashion to the cost imputations. As the period fixed effects are common to all airlines and airports, the only two elements left to estimate are the route-airline fixed effects and the demand shocks. For the demand shocks, I take random draws from the demand estimation residuals. On the other hand, I decompose the route-airline fixed effects as follows:

$$\eta_{imr} = \theta_{d1} \cdot dist_r + \theta_{d2} \cdot dist_r^2 + \vartheta_{ia} + \varsigma_{ia'} + \varepsilon_{imr}.$$

In words, I decompose the demand between two airport-airline specific fixed effects and a distance effect. When ϑ_{ia} or $\varsigma_{ia'}$ do not exist for some airline-airport combination, I assume that parameter's value is equal to the lowest available value for that airport. As before, this aims to capture the fact that if an airline is not flying to an airport, consumers in that airport should not value that airline more than an airline that does fly there. Finally, I construct the market size for the missing markets in a similar way to what I did for the existing markets.

I impose additional constraints to the problem to avoid reaching unrealistic results. Specifically, I impose the restriction that airlines cannot fly a total number of flights that is higher than the amount of flights they fly in the data. In addition, I include the constraint that each flight can only handle up to 1.25 times the number of passengers per flight that I see in the data—which matches the load factor observed in the industry.[40] Finally, I impose the restriction that for each base they have, airlines cannot operate more than 1.25 times the number of flights they operate in the data. This aims to capture the fact that some operational factors—like hangars, crews, mechanics—are tied to certain bases, which means that airlines cannot freely move their whole operation to a different base. These constraints are needed to address operational limits that my model does not take into

[40] See, for example, IATA (2025).

account for the sake of simplicity.

Given Assumption 5, I simulate a static version of the Model in Section 4, with the additional restriction that airlines cannot fly less than 80% of their slots at slot constrained airports. I solve this equilibrium through iterated best response for 2016. With Assumption 5 the simulation gives me a stylized estimate for the long-run results, as the assumption implies slots never change hands. While we know that for a given frequency profile, the ensuing price choice game has a unique equilibrium, the same cannot be said for the frequency game. While I run my iterated best-response algorithm from various starting points to address this concern, I cannot rule out that the data come from one of many equilibria.[41]

In this section, I compare my baseline simulation against the data. I summarize these differences in Table 5. I find that even with the additional constraints I impose, my simulation still overestimates the airline profits with respect to the actual data. In addition, I find that consumer surplus is lower than in the data, which leads to an overall net negative difference in total welfare. Finally, I find that the total number of markets is 20% lower than in the data. Overall, around 62% of all markets active in the baseline are also active in the data.

	Baseline	**Data**	**Change**
Airline Profit	\$12.7 Bn	\$7.3 Bn	73.31%
Consumer Surplus	\$14.6 Bn	\$21.1 Bn	-30.87%
Total Welfare	\$27.3 Bn	\$28.5 Bn	-4.00%
Passengers	504.9 Mn	532.9 Mn	-5.24%
Average Markup	36.05%	36.54%	1.37%
Daily Flights	16,824	17,131	-1.79%
Total Markets	4,772	6,014	-20.66%

Table 5: Comparison between the baseline simulation—where airlines are required to fly at least 80% at each slot controlled airport—and the data. In both cases, the slot allocation is the actual slot allocation of 2016.

Furthermore, in Figure 9 we can see how well the baseline's chosen prices and frequencies match with the data on the routes active in both the data and the baseline. We see that there is a very good match for the prices, but not for the frequencies. In general, airlines overfly the existing routes in the simulation. As I mention in Section 6.2, my marginal cost of adding one additional flight may be underestimated, specially for the busiest routes. Given that these routes usually include slot-controlled airports, this bias induces airlines to fly more to these airports, irrespective of whether they have slot controls or not. Thus, the bias makes the baseline and counterfactual more similar, which means that my results are a lower bound of the actual effects of grandfathering on airline behavior.

There are additional operational factors I do not include in my model, like the flight time or the specific locations aircraft need to be to serve certain routes. The bias of this second factor on my

[41] See Appendix G for a description of the randomization procedure and a summary of the algorithm's result from various seeds.

results is unclear, as it depends on how easily accessible the most desirable airports are. Still, given that these desirable airports are probably slot constrained, this would also bias airlines to fly more to slot constrained airports.

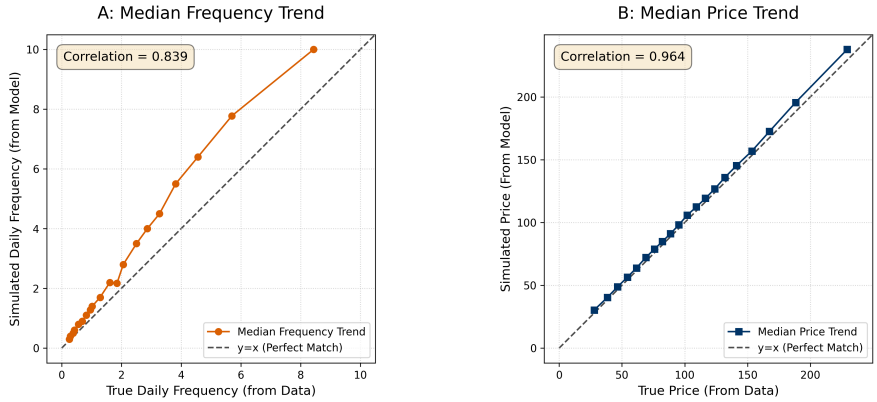


Figure 9: Binscatter comparing simulated and actual prices for routes operated in both the Baseline simulation and the data. The 45 degree line shows a perfect match. We see that prices match very well with the data, as the marginal cost of an extra seat has no bias in its estimation. On the other hand, the marginal cost of an extra flight is biased downwards for the largest airports, which explains the worse match between data and simulation for frequencies. Still, this bias makes my baseline more similar to the counterfactual simulations, as airlines have an incentive to fly more regardless of the 80% requirement.

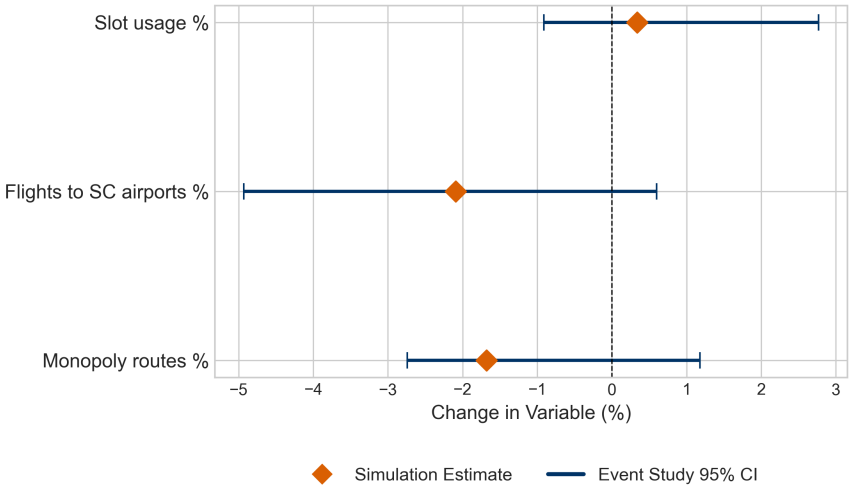


Figure 10: Match between average event study confidence intervals in Section 3 and simulated change from introducing slot controls at the 8 airports from the event study. I run my simulation for each quarter of the first year where each airport becomes slot controlled under two scenarios—one where the airport has slot controls, and another one where it does not have controls. I then compute the mean of the 3 event study variables for both cases. All 3 differences of means fall within the event study's average confidence intervals, indicating a good match between the event study and my model.

To check how my model matches changes in grandfathering, I check how well it predicts the airline behavior I see in my event study. To do this, I run the simulation for all quarters where one of the 8 airports from the event study in Section 3 changed from not having slot controls to having them. I run each quarter twice, once with the true slot controlled airports for that quarter, and once assuming that the airport that just switched to having slot controls still has no controls.[42] For each simulation, I store the route-airline tuples for the airports that started having slot controls in that quarter.

I then measure the average change in the share of monopoly routes, flights to other slot controlled airports and slot usage across all 8 airports for their relevant quarters. In Figure 10 I compare the average change in the three variables with the average post-treatment confidence intervals for the three event studies in Section 3. We see that all 3 simulation changes fall within the event study confidence intervals. This shows that while my baseline has some clear biases with respect to the actual data, it accurately matches the changes in the data that grandfathering produces.

8.3 Counterfactual I: No 80% usage requirement with Current Slot Allocation

I solve the system-wide equilibrium for the Model in Section 4 for 2016 assuming airlines only care about stage profits and Assumption 5 is not in place. This corresponds to the case where the 80% usage rule is not in effect, meaning that the airlines own the slots in perpetuity without any usage requirements. While this counterfactual does not correspond to any current reform proposal, it allows me to understand the pro-competitive effects of the rule without confounding them with the potential efficiency gains from reallocating airport slots through an auction.

	Baseline	**Counterfactual I**	**Change**
Airline Profit	\$12.7 Bn	\$13.5 Bn	6.36%
Consumer Surplus	\$14.6 Bn	\$12.4 Bn	-14.85%
Total Welfare	\$27.3 Bn	\$26.0 Bn	-4.97%
Passengers	504.9 Mn	475.3 Mn	-5.86%
Average Markup	36.05%	36.71%	1.84%
Daily Flights	16,824	16,553	-1.61%

Table 6: Comparison between the baseline simulation in Section 8.1—where airlines are required to fly at least 80% at each slot controlled airport—and the first counterfactual simulation—where the 80% usage requirement doesn't apply. In both cases, the slot allocation is the actual slot allocation of 2016.

[42] For example, for quarter 2 and 3 of 2016, I run the regression with Hamburg Airport having and not having slot controls, as Hamburg airport started having slot controls during Summer of 2016.

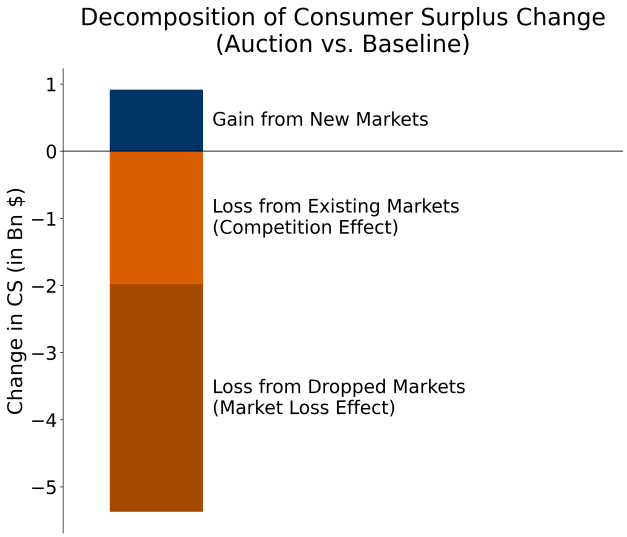


Figure 11: Consumer welfare change decomposition. While the net effect is negative, we see that the drop in consumer welfare comes from two sources: first, airlines stop serving some markets without the 80% requirement, which I call the Market Loss Effect; second, airlines increase their market power on other markets, which I call Competition Effect. Although consumers in new markets are better off, this does not compensate the losses from existing markets.

Table 6 shows some key summary statistics comparing the baseline from Section 8.2 and the computed equilibrium under no usage requirement. We can see that eliminating the usage requirement results in a large shift in welfare from consumers to airlines. This happens due to two main reasons. First, airlines compete in fewer markets as a result of the rule change, which leads to a 1.84% increase in average markup. Second, airlines serve fewer markets, which also accounts for the

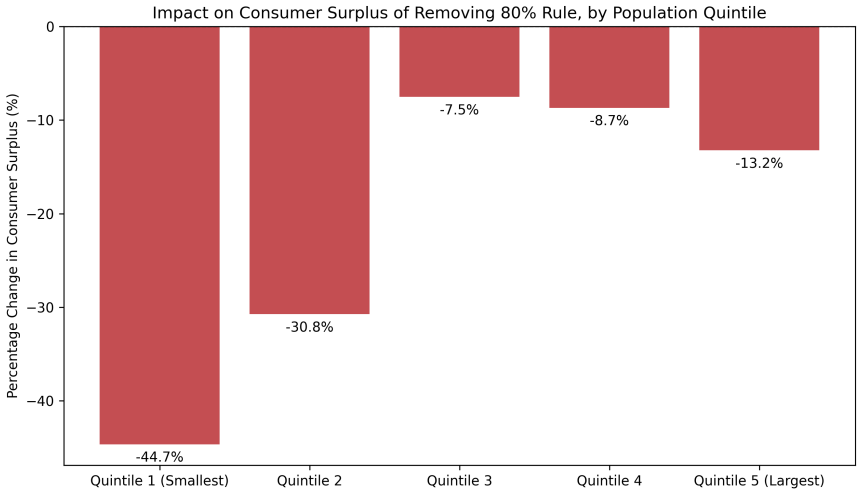


Figure 12: Change in consumer welfare by quintile after switching from the baseline with the 80% usage requirement to the counterfactual with no usage requirement. The first quintile includes all catchment areas with a population below 67 thousand inhabitants. The fifth quintile include all catchment areas with a population greater than 1.1 million inhabitants.

drop in overall consumer welfare. I decompose the total welfare loss between the loss due to lower competition and to dropped markets in Figure 11. In addition, I show in Figure 12 how eliminating grandfathering affects different cities. We see that smallest cities (with less than 67,000 inhabitants) are the hardest hit, which falls in line with the large market loss effect.

My simulation also allows me to study how the use-it-or-lose-it rule affects congestion. Intuitively, if grandfathering causes airlines to fly more between slot constrained airports, then there may be airports that are congested solely *because* they have slot controls. While I do find that some airports see large drops in their usage, this by itself does not guarantee that there would be no more congestion. This is because my model takes non-Schengen area flights as given. To address this issue, I take a conservative approach and declare an airport as uncongested for a season in my simulation only if it sees a drop of at least 33% in total Schengen Flights, and if Schengen Flights represent at least 60% of total airport flights.

Airport (IATA code)	Season	Flight Change	Schengen %
Barcelona (BCN)	Winter	-38.5%	76.5%
Bergen (BGO)	Winter	-43.6%	94.2%
Berlin (TXL)	Winter	-37.5%	83.9%
Catania (CTA)	Winter	-37.7%	95.2%
Corfu (CFU)	Summer	-42.0%	60.1%
Genoa (GOA)	Summer	-34.1%	83.8%
Gran Canaria (LPA)	Winter	-34.9%	85.7%
Madrid (MAD)	Winter	-40.3%	71.8%
Menorca (MEH)	Winter	-45.3%	70.2%
Milan (BGY)	Summer	-55.3%	78.3%
Milan (BGY)	Winter	-50.2%	86.8%
Mykonos (JMK)	Summer	-61.1%	84.1%
Paris (ORY)	Winter	-34.9%	77.1%
Rome (CIA)	Summer	-33.8%	64.8%
Rome (CIA)	Winter	-77.5%	67.1%
Santorini (JTR)	Summer	-52.4%	88.1%

Table 7: List of airport-season pairs that would become uncongested without the grandfathering rule. I define an airport as uncongested if Schengen Area flights make up more than 60% of total airport traffic, and if the simulation sees a reduction of more than 33% in total Schengen area flights.

Table 7 presents the list of uncongested airport-season pairs. I find that 14 airports, or 20% of all slot constrained airports, have at least one season where there would be no congestion without the grandfathering rule. In general, these are small to medium size airports, but the list also includes the Winter seasons of 3 of the top 25 busiest airports in Europe (Barcelona, Madrid and Paris-Orly). In addition, the list also includes the low-cost airports for Rome (Ciampino) and Milan (Bergamo).

The total amount of passengers drops slightly (-6%). Assuming that the average intra-Schengen passenger flies 2 hours, and if we assume that the average CO_2 emissions per passenger for such a flight are 85 Kg, we can see that the total CO_2 emissions of intra-Schengen flights drops by around 2.5 million metric tons per year. This represents around 0.07% of yearly EU CO_2 emissions. In a recent meta-analysis, the US Environmental Protection Agency (EPA) reports a 95%-confidence

interval for the Social Cost of Carbon (SCC) of between $33 and $285 per metric ton of CO_2.[43] Thus, the lower emissions due to the elimination of the 80% usage requirement would result in a welfare gain of between $82.5 and $712.5 million per year. Nevertheless, we can see that even the higher end of this estimate would still result in an overall welfare loss.

8.4 Counterfactual II: No 80% Usage Rule & Efficient Slot Reallocation

In this section I reallocate airport slots to find the allocation that maximizes total airline profits without the 80% usage rule rule. This is an approximation to the outcome if slots were assigned using a combinatorial auction. Rassenti et al. (1982) show that such an auction maximizes total bidder surplus, which in my context is the sum of total (current period) profits. I use a genetic algorithm to bypass the curse of dimensionality in the number of possible allocations. I find that in this scenario, total welfare is still lower than the baseline scenario, and while airlines would earn higher gross profits, their profits net of payments would still be lower than the baseline.

Given that in 2016 there were 245,027 airport slots in my dataset, a brute force query through all possible slot combinations is not computationally feasible. In addition, the many overlapping slot constraints mean that using standard optimization tools is also impracticable. Thus, in this section I develop a genetic algorithm to find the slot allocation that maximizes total airline profits.[44] Recent papers have shown that this approach can yield efficiency gains in the context of airport slot allocation (Melder et al., 2023; Katsigiannis and Zografos, 2024). My genetic algorithm searches over the space of all possible slot allocations that respect the total slot capacity at each constrained airport. The process begins with an initial population of 64 distinct slot allocations. To ensure the optimized outcome is at least as good as Section 8.3's simulation, one of these initial allocations is the observed real-world allocation, while the remaining 63 are generated by randomly reallocating the 245,027 Schengen-area slots among airlines.

For each proposed allocation, the algorithm's solves for the full Static Equilibrium from Section 8.3, and calculates the resulting total airline profits, which will be the fitness function the algorithm aims to maximize. The algorithm then iteratively improves the population over 100 generations using three evolutionary operators. First, I run 64 tournaments where I compare 3 random allocations and choose the fittest one to become a parent for the next generation. Second, I use a single-point crossover operator where I swap one airport's whole slot allocation between two parents to create a new slot allocation. Finally, I also include a mutation operator that randomly reallocates 10,000 slots between airlines. For each tournament winner, the algorithm applies the crossover operator with probability 0.7, and the mutation operator with probability 0.4. I store the best-performing allocation ever found and I take this to be an approximation the outcome of an efficient slot auction. Figure 13 presents an example of the algorithm's convergence for the second quarter of 2016.

To approximate the payments under a slot auction, I recompute the genetic algorithm excluding one airline at a time from the slot reallocation. This allows me to construct the VCG payments that

[43] These are the values assuming a Ramsey discount rate of 2.5%. For more details, see EPA (2023).

[44] While this approach is guaranteed to reach the optimum in the limit (Altenberg, 1995), in practice computational constraints only allow me to find an improvement over the current allocation.

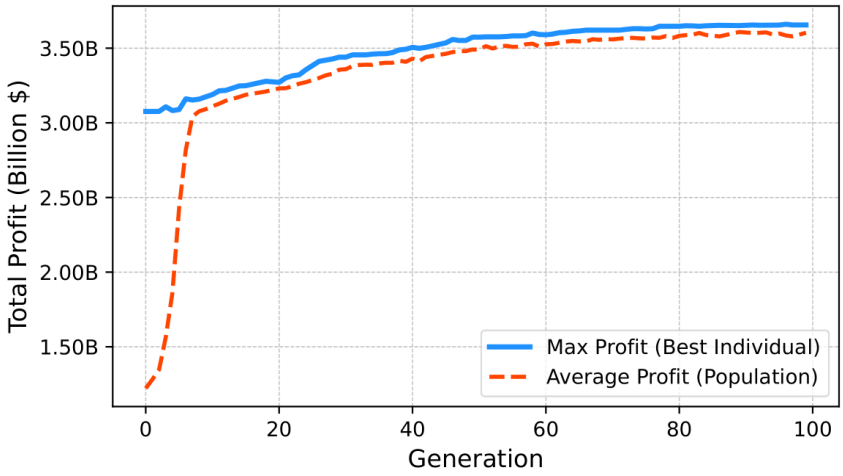


Figure 13: Algorithm Convergence Example for the second quarter of 2016. The genetic algorithm starts with a population of 64 individuals, and over 100 generations it selects the fittest individuals—i.e., the slot allocations with the highest total airline profits. Random mutations and selective crossover between allocations help the algorithm to avoid local maxima.

each airline would have to make under an efficient auction.[45] Formally, the payment that airline i has to make is given by:

$$payment_i = \sum_{j \neq i} \left(\Pi_j^*(\text{i cannot get slots}) - \Pi_j^*(\text{i can get slots}) \right). \tag{9}$$

Given that the optimization process grows exponentially harder with each additional airline, I run these VCG algorithms for only 25 generations, to make sure the solutions of both algorithms are comparable.

Table 8 gives some summary statistics for the algorithm's optimal allocation, comparing it to the baseline simulation from Section 8.1. We see that although gross airline profits are higher under the auction, consumers are worse off as a result of the auction. As before, this happens both because airlines exit from markets to concentrate their operations on particularly profitable routes, as well as due to the lower output and competition levels. Indeed, we see that the average markup is even higher than on the first counterfactual. As a result, overall welfare is lower under this auction counterfactual than under the baseline scenario. This shows that the pro-competitive mechanism of grandfathering dominates the efficiency gains from regularly auctioning airport slots. In addition, we see that the drop in total flights is particularly high for routes between slot constrained (SC) airports.

[45] A required assumption when using VCG payments is that the excluded player's payoff is independent of the auction winner's identity. This may not be true for slot auctions due to the network effects—for example, AirFrance-KLM winning a slot at Frankfurt airport would hurt Lufthansa much more than Air Malta winning that same slot. The fact that a large majority of airports in my sample (and indeed most airline bases) are slot-controlled airports means that excluding an airline from acquiring slots *at all airports* greatly limits their choices. Thus, the excluded airline's payoff will be approximately the same regardless of the resulting slot allocation among all other airlines. I expand on this in Appendix H.

	Baseline	Counterfactual II	Change
Airline Profit (gross of auction payments)	\$12.7 Bn	\$15.4 Bn	20.63%
Airline Profit (net of auction payments)	\$12.7 Bn	\$9.6 Bn	-23.16%
Consumer Surplus	\$14.6 Bn	\$11.7 Bn	-19.65%
Total Welfare	\$27.3 Bn	\$27.0 Bn	-1.11%
Passengers	504.9 Mn	434.5 Mn	-13.93%
Average Markup	36.05%	46.56%	44.31%
Daily Flights	16,824	14,788	-12.10%
Daily Flights, SC to SC	8,055	6,948	-13.70%
Daily Flights, SC to Non-SC	7,580	6,750	-11.00%
Daily Flights, Non-SC to Non-SC	1,189	1,091	-8.02%

Table 8: Comparison between the baseline simulation in Section 8.1—where airlines are required to fly at least 80% at each slot controlled airport—and the second counterfactual simulation—where the 80% usage requirement doesn't apply. In the baseline, the slot allocation is the real slot allocation of 2016. In the counterfactual, it is the slot allocation that maximizes total airline profits according to this section's genetic algorithm. This corresponds to an approximation of the allocation under an efficient auction. The auction payments are the VCG payments from Equation (9). SC refers to a slot-constrained airport.

This shows that airlines tend to fly between slot controlled airports to meet usage requirements at origin and destination with a single flight.

Figure 14 decomposes the welfare effect of running an auction across Europe NUTS 1 regions.[46] This map shows two facts. First, most regions are worse off in the auction counterfactual. This is a direct consequence of the fact that most regions have a Level 3 airport that benefited from the 80% usage requirement. Second, we see that a few regions (Like North Norway, Brittany or Slovenia) are better off after the auction. This happens because these areas' main airports are not Level 3. While airlines would find it profitable to fly to these regions, the 80% usage requirement induces them to prioritize the continuation value of their slot portfolio.

Under this counterfactual, there is a reduction of 6.0 million metric tons of CO_2 due to the lower passenger numbers. Using the EPA 95%-confidence interval for the SCC that I used in Section 8.3, I find that this drop in CO_2 emissions results in a welfare gain of between \$197 and \$1,704 million per year. While the lower end of this estimate still does not cover the welfare loss due to the lower competition, the higher end results in a 5% increase in total welfare. However, consumers are still worse off under the auction even if they accrue the benefits from the lower emissions.

While gross airline profits are higher under the auction, airlines are actually worse off once payments are taken into account. This helps explain why airlines are overwhelmingly against airport slot auctions (IATA, 2024). Table 9 disaggregates the results across all airlines. We see that most airlines are worse off as a result of the auction, and that the relative slot ownership does not

[46] For Hungary, Netherlands and Belgium, I use the NUTS 0 regions, as in my model each of these countries is its own catchment area. For Norway, Portugal and Finland, I use the NUTS 2 regions with some level of aggregation across regions, as the NUTS 1 regions are too aggregated given my model. I also aggregate some NUTS 1 regions for Poland, France and Germany.

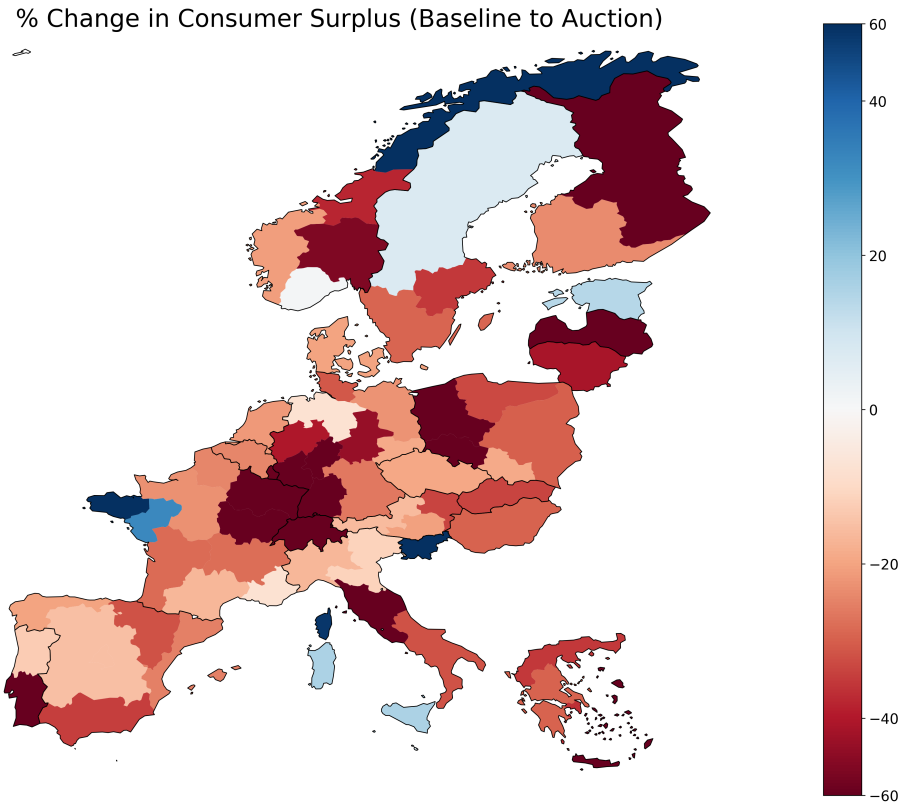


Figure 14: Schengen Area NUTS 1 regions by change in consumer welfare when going from the baseline—where the 80% requirement is in place—to the auction counterfactual—where the 80% requirement is not in place and slots are reallocated to maximize total gross airline profits. We see that most regions are worse off due to the auction, as without the 80% requirement airlines fly fewer flights. This results in a drop in competition and offered varieties. Some regions with low population densities see an increase in welfare.

change too much across airlines. The largest lowcost carrier, Ryanair, is better off even after taking into account auction payments. Although no airlines have expressed support for implementing slot auctions, this outcome makes sense given that Ryanair has been particularly vocal about its disconformity with the current slot allocation system.[47]

The total auction revenue generated—around $ 5.3 Bn per year—falls in line with the revenue from other essential input auctions, notably spectrum (Milgrom, 2021). In this paper, I do not make any assumptions on what that revenue would be used for. One possibility would be to allocate that revenue for airport expansion at particularly congested airports. If this was the case, the total welfare might actually be positive over the long run, although there still might exist distributional effects between airlines and/or cities.

On the other hand, my outcome approximates an auction where all actors can perfectly compute their profits for each possible slot configuration. Given the sheer scale of the European slot portfolio, this is a strong assumption. Still, it presents a "best case scenario" for the efficiency gains due to

[47]See, for example Ryanair (2022).

	Slots		Profits (Mn $)			
				Auction		
Airline	Baseline	Auction	Baseline	Gross	Payments	Net
Ryanair	17,281	16,343	2,670.06	3,823.49	479.83	3,343.66
Lufthansa Group	55,998	49,215	2,118.10	1,730.06	484.91	1,245.15
EasyJet	15,770	15,209	1,333.69	1,348.89	148.58	1,200.31
Air France-KLM	36,703	35,359	1,308.91	1,531.72	574.18	957.55
IAG (BA-Iberia)	26,471	24,547	1,234.73	1,267.42	364.38	903.04
Air Berlin	14,846	14,342	961.73	1,223.12	845.06	378.06
Norwegian	13,170	13,035	911.34	803.36	535.49	267.87
SAS Group	19,484	18,533	711.95	937.31	749.32	187.99
Air Europa	2,570	2,742	330.20	327.77	93.40	234.37
Wizz Air	1,456	1,599	322.02	368.59	175.83	192.76
Aegean Airlines	5,803	6,461	203.60	98.44	47.13	51.31
Finnair	4,731	5,191	193.29	177.77	113.13	64.64
Alitalia	12,578	12,871	160.64	340.18	206.11	134.08
Volotea	1,871	2,079	155.50	119.55	59.25	60.30
LOT Polish Airlines	4,099	4,660	108.50	99.24	35.73	63.50
Czech Airlines	1,726	2,120	75.47	97.58	66.82	30.76
Air Malta	430	485	55.66	44.30	39.80	4.50
Luxair	971	1,066	50.07	59.98	40.48	19.51

Table 9: Profit and Slot Ownership disaggregation under the Baseline scenario of Section 8.1 and this section's auction. I compute the VCG payments according to Equation (9).

the auction format. As the real outcome would probably be worse than this, my estimates on the total welfare loss due to auctions probably represent a lower bound to the actual losses.

9 Conclusion

In this paper I study the effects of the grandfathering rule on airline competition and welfare in the European Union. I show that the 80% usage requirement, despite contributing to congestion, generates powerful pro-competitive effects that deliver immense value to consumers. By inducing airlines to schedule flights between slot-controlled airports simply to retain future access, the rule increases total flights and intensifies competition. I estimate that severing the link between this year's slot usage and next year's slot allocation would lead to an overall loss in welfare. In addition, I construct a genetic algorithm to approximate the outcome of a slot auction. While there are efficiency gains from this slot reallocation, they are not high enough to compensate the lower competition levels due to the drop in route commonality.

I develop and estimate a dynamic structural model of airline network choice and price competition to quantify the system-wide impact of the rule. I bypass the curse of dimensionality that my model has by assuming airlines behave as if they were required to use their slots at least 80% of the time at each slot controlled airport. This is a stylized simplification of airline behavior that matches reduced-form evidence in my paper. I compare this baseline scenario to the equilibrium without

usage requirements, holding existing the slot allocation fixed. This allows me to isolate the current system's pro-competitive effects. I show that airlines would reconfigure their networks without the rule, which would result in a 15% drop in consumer welfare, and a 5% drop in total welfare. Paradoxically, this would de-congest 20% of currently constrained airports.

Next, I simulate the outcome of regular system-wide slot auctions, with the objective of understanding the trade-off between allocative efficiency and lower competition. To do this, I find the slot allocation that maximizes total airline profit, which is the outcome of a standard combinatorial auction (Rassenti et al., 1982). As my dataset has over 246,000 slots, and given the large number of intertwined slot constraints, I cannot employ standard optimization techniques or brute force to find this profit maximizing allocation. Instead, I employ a genetic algorithm where I evolve a set of random slot allocation over 80 generations, selecting for those that produce the highest total profit levels.

I find that the auction leads to a 1% drop in total welfare with respect to the baseline. Thus, while reallocating airport slots results in some efficiency gains, eliminating the link between today's usage and tomorrow's slot allocation still leads to an overall drop in total welfare due to the lower competition. In particular, I find that consumer welfare drops even further after the auction. Furthermore, even though there is a reduction in total CO_2 emissions due to the drop in total passengers, even under the EPA's largest social cost of carbon estimates consumers are still worse off under this auction counterfactual. On the other hand, I use my genetic algorithm to compute an approximation to the VCG payments airlines would make under an efficient auction. I find that even though total profits are higher under an auction, airlines are worse off after taking payments into account. This may explain why airlines generally oppose auctions (IATA, 2024).

My results highlight a critical, previously unquantified trade-off for policymakers and they open several avenues for future research. First, while my analysis models the complete elimination of the rule, future work could explore the marginal effects of alternative usage thresholds. A lower threshold, for instance, might offer a better balance between fostering competition and allowing airlines network flexibility. Second, my framework could be used to study the introduction of a secondary trading market, a policy proposal currently under consideration by the European Commission. Finally, my paper's framework could be extended to other essential inputs markets, such as electromagnetic spectrum or water and land rights.

References

Victor Aguirregabiria and Chun-Yu Ho. A dynamic game of airline network competition: Hub-and-spoke networks and entry deterrence. *International Journal of Industrial Organization*, 28 (4):377–382, 2010.

Victor Aguirregabiria and Chun-Yu Ho. A Dynamic Oligopoly Game of the US Airline Industry: Estimation and Policy Experiments. *Journal of Econometrics*, 168(1):156–173, 2012.

Airport Coordination Limited. Completed slot trades, 2025. URL https://www.acl-uk.org/completed-slot-trades/. Accessed on July 30, 2025.

Lee Altenberg. The schema theorem and price's theorem. In *Foundations of genetic algorithms*, volume 3, pages 23–49. Elsevier, 1995.

Isaiah Andrews, James H Stock, and Liyang Sun. Weak instruments in instrumental variables regression: Theory and practice. *Annual Review of Economics*, 11(1):727–753, 2019.

Gaurab Aryal, Charles Murry, and Jonathan W Williams. Price discrimination in international airline markets. *Review of Economic Studies*, 91(2):641–689, 2024.

Australian Government. Reforms to the Sydney Airport demand management framework, 2024. URL https://www.infrastructure.gov.au/infrastructure-transport-vehicles/aviation/airports/reforms-sydney-airport-demand-management-framework. Accessed on May 29, 2025.

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

Michael O Ball, Lawrence M Ausubel, Frank Berardino, Peter Cramton, George Donohue, Mark Hansen, and Karla Hoffman. Market-based alternatives for managing congestion at New York's LaGuardia airport. Citeseer, 2007.

Sebastián D Bauer. Aftermarket deterrence in essential input auctions. 2025.

C Lanier Benkard, Aaron Bodoh-Creed, and John Lazarev. Simulating the dynamic effects of horizontal mergers: US airlines. *Manuscript, Yale University*, 2010.

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

Steven Berry, Michael Carnall, and Pablo T Spiller. Airline hubs: costs, markups and the implications of customer heterogeneity. *Competition policy and antitrust*, 2006.

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

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

Steven T Berry, Michael Carnall, and Pablo T Spiller. Airline hubs: costs, markups and the implications of customer heterogeneity, 1996.

Jose M Betancourt, Ali Hortaçsu, Aniko Oery, and Kevin R Williams. Dynamic price competition: Theory and evidence from airline markets. Technical report, National Bureau of Economic Research, 2022.

Martin Bichler, Peter Cramton, Peter Gritzmann, and Axel Ockenfels. It is time to auction slots at congested airports. *ECONtribute Policy Brief*, 2021a.

Martin Bichler, Richard Littmann, and Stefan Waldherr. Trading airport time slots: Market design with complex constraints. *Transportation Research Part B: Methodological*, 145:118–133, 2021b.

Martin Bichler, Peter Gritzmann, Paul Karaenke, and Michael Ritter. On Airport Time Slot Auctions: A Market Design Complying with the IATA Scheduling Guidelines. *Transportation Science*, 57(1):27–51, 2023.

Boeing. Commercial Market Outlook 2024–2043, 2024.

Christian Bontemps, Kevin Remmy, and Jiangyu Wei. Ex-post evaluation of the American Airlines–US Airways merger: a structural approach. *Journal of Transport Economics and Policy (JTEP)*, 56(2):129–155, 2022.

Christian Bontemps, Cristina Gualdani, and Kevin Remmy. Price competition and endogenous product choice in networks: Evidence from the us airline industry. Technical report, Working Paper, 2023.

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

Severin Borenstein and Nancy L Rose. Competition and price dispersion in the US airline industry. *Journal of Political Economy*, 102(4):653–683, 1994.

Kirill Borusyak, Xavier Jaravel, and Jann Spiess. Revisiting event-study designs: robust and efficient estimation. *Review of Economic Studies*, page rdae007, 2024.

Federico Ciliberto and Elie Tamer. Market structure and multiple equilibria in airline markets. *Econometrica*, 77(6):1791–1828, 2009.

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

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

Edward H Clarke. Multipart pricing of public goods. *Public choice*, pages 17–33, 1971.

Daniele Condorelli. Efficient and equitable airport slot allocation. *Rivista di politica economica*, 1 (2):81–104, 2007.

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

Rebekah Dix and Roi Orzach. Market power spillovers across airline routes. Technical report, 2022.

EPA. EPA Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances. *Washington, DC: US IAWG*, 2023.

European Commission. Allocation of EU airport slots - review of rules, 2023. URL https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/13528-Allocation-of-EU-airport-slots-review-of-rules_en. Accessed on May 29, 2025.

Billy Ferguson. Trade frictions in surface water markets. Technical report, Working paper, 2024.

Xiaowen Fu, Zheng Lei, Shaoxuan Liu, Kun Wang, and Jia Yan. On-time performance policy in the chinese aviation market-an innovation or disruption? *Transport Policy*, 95:A14–A23, 2020.

Ian Gale and Daniel P O'Brien. The welfare effects of use-or-lose provisions in markets with dominant firms. *American Economic Journal: Microeconomics*, 5(1):175–193, 2013.

GAO. Slot-controlled airports: FAA's rules could be improved to enhance competition and use of available capacity. *Report to the Committee on Commerce, Science, and Transportation, U.S. Senate*, pages 12–902, 2012.

GAO. Spectrum management: FCC's use and enforcement of buildout requirements. *Report to Congressional Requesters*, pages 1–70, 2014.

Government of the United Kingdom. Airport slot allocation system reform, 2024. URL https://www.gov.uk/government/consultations/airport-slot-allocation-system-reform. Accessed on May 29, 2025.

Theodore Groves. Incentives in teams. *Econometrica*, pages 617–631, 1973.

IAG. *Full Year Results 2024*. 2024.

IATA. *Air/Rail Inter-Modality Study. Final Report*. 2003.

IATA. Estimating air travel demand elasticities. 2007.

IATA. Slots white paper: Future airport slot policy and the airline industry, 2024.

IATA. Worldwide Airport Slot Guidelines, 2025.

IATA. *Quarterly Air Transport Chartbook*. 2025.

Maarten CW Janssen and Vladimir A Karamychev. Do auctions select efficient firms? *The Economic Journal*, 120(549):1319–1344, 2010.

Bernhard Kasberger. When can auctions maximize post-auction welfare? *International Journal of Industrial Organization*, 89:102972, 2023.

Fotios Katsigiannis and Konstantinos Zografos. A multi-objective genetic algorithm for airport slot allocation decision-making. *Available at SSRN 4742975*, 2024.

Frank Kleibergen and Richard Paap. Generalized reduced rank tests using the singular value decomposition. *Journal of econometrics*, 133(1):97–126, 2006.

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

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

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

Jeffrey B Liebman and Neale Mahoney. Do expiring budgets lead to wasteful year-end spending? evidence from federal procurement. *American Economic Review*, 107(11):3510–3549, 2017.

Lufthansa Group. *2024 Annual Report*. 2024.

Marleen Marra. A Market for Airport Slots. *Available at SSRN: https://ssrn.com/abstract=4948006 or http://dx.doi.org/10.2139/ssrn.4948006*, 2024.

David Melder, John Drake, and Sha Wang. An evolutionary hyper-heuristic for airport slot allocation. In *International Conference on the Applications of Evolutionary Computation (Part of EvoStar)*, pages 53–68. Springer, 2023.

Paul Milgrom. Auction research evolving: Theorems and market designs. *American Economic Review*, 111(5):1383–1405, 2021.

Toshihide Mizuno. On the existence of a unique price equilibrium for models of product differentiation. *International Journal of Industrial Organization*, 21(6):761–793, 2003.

Volker Nocke and Nicolas Schutz. Multiproduct-firm oligopoly: An aggregative games approach. *Econometrica*, 86(2):523–557, 2018.

Angela Nuic, Damir Poles, and Vincent Mouillet. Bada: An advanced aircraft performance model for present and future atm systems. *International journal of adaptive control and signal processing*, 24(10):850–866, 2010.

Amedeo R Odoni. A review of certain aspects of the slot allocation process at level 3 airports under regulation 95/93. 2021.

Lale Özbakir, Adil Baykasoğlu, and Pınar Tapkan. Bees algorithm for generalized assignment problem. *Applied Mathematics and Computation*, 215(11):3782–3795, 2010.

Stephen J Rassenti, Vernon L Smith, and Robert L Bulfin. A Combinatorial Auction Mechanism for Airport Time Slot Allocation. *The Bell Journal of Economics*, pages 402–417, 1982.

Itzchak Tzachi Raz. Use it or lose it: Adverse possession and economic development. *Unpublished Manuscript*, 2018.

Peter C Reiss and Pablo T Spiller. Competition and entry in small airline markets. *The Journal of Law and Economics*, 32(2, Part 2):S179–S202, 1989.

James D Reitzes, Brendan McVeigh, Nicholas Powers, and Samuel Moy. Competitive effects of exchanges or sales of airport landing slots. *Review of Industrial Organization*, 46(2):95–125, 2015.

Patrick Rey and David J Salant. Auctions for essential inputs. *Available at SSRN 2941357*, 2017.

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

Ryanair. Ryanair Solves Lufthansa's "Ghost Flight" Problem – Just Sell The Seats To Consumers At Low Fares!!!, 2022. URL https://corporate.ryanair.com/news/ryanair-solves-lufthansas-ghost-flight-problem-just-sell-the-seats-to-consumers-at-low-fares/ Accessed on May 29, 2025.

Schedule Coordination Austria. Slot monitoring and historic rights, 2024. URL https://www.slots-austria.com/jart/prj3/sca_2024/uploads/data-uploads/Slot%20Monitoring%20and%20historic%20rights%20V4.pdf. Accessed on July 30, 2025.

James Schummer and Azar Abizada. Incentives in landing slot problems. *Journal of Economic Theory*, 170:29–55, 2017.

James Schummer and Rakesh V Vohra. Assignment of arrival slots. *American Economic Journal: Microeconomics*, 5(2):164–185, 2013.

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

Pere Suau-Sanchez, Guillaume Burghouwt, and Montserrat Pallares-Barbera. An appraisal of the CORINE land cover database in airport catchment area analysis using a GIS approach. *Journal of Air Transport Management*, 34:12–16, 2014.

The Economist. Why we're advocating the case for change in airline slots, 2017. URL https://www.economist.com/leaders/2017/11/16/the-system-for-allocating-airport-slots-is-broken. Accessed on Sep 25, 2025.

William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.

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

Zhe Yuan and Panle Jia-Barwick. Network competition in the airline industry: An empirical framework. Technical report, National Bureau of Economic Research, 2024.

Appendix A Adding Hour-Day Valuations

In this section I show that the problem in Section 4 can be rewritten to allow for hour-day specific prices and tastes. Under Assumptions 2-4, the final demand will look identical to the demand in the main text. These are strong behavioral assumption on the firm, and they could be lifted with flight-level data.

The indirect utility of a traveler that purchases flight hd by airline i for route $r = a \to a'$ in market m during period t will now be:

$$U(p^{hd}_{imrt}, \nu^{hd}_{imrt}) = \overbrace{\beta_{p,m} \cdot p^{hd}_{imrt}}^{\text{price}} + \overbrace{\beta^{hd}_{imrt}}^{\text{hour/day dummy}} + \overbrace{X_{imrt} \cdot \beta_x}^{\text{Other obs charac.}} + \overbrace{\xi_{imrt}}^{\text{Unobs. airline-market charac.}} + \overbrace{\nu^{hd}_{imrt}}^{\sim iid\ T1EV}.$$

Here, p^{hd}_{imrt} is the average price that airline i charges for flight hd in route r in market m during period t. Consumers now derive different utilities depending on the flight's departure hour and date, and the effect of these different hour-date combinations is given by β^{hd}_{imrt}. Notice that I allow for hour-date taste to vary between markets, routes, airlines and periods. All other elements are analogous to the main text.

Notice that the only two things that vary by flight in this version of the model are the price and the departure time and weekday. This assumption is essential to aggregate the model up to the airline-route level. As most other flight-level characteristics (like crew-experience, aircraft-type, meal service, to mention a few) are either unobservable at the time of booking, or do not vary significantly between two flights by the same airline to the same destination, the assumption is not a strong one to make.

Appendix A.1 Airline Problem

Each airline now has two teams responsible of pricing. First, there is an airline-wide team, responsible of choosing a route-level price component p_{imrt} for each route r in market m where airline i chose to enter during period t. Next, each airline i has a day-to-day operations team for each market-route tuple. The day-to-day operations team sees p_{imrt} and has to then choose an hour-day specific component, μ^{hd}_{imrt}. Thus, the price that a consumer sees can be rewritten as:

$$p^{hd}_{imrt} = p_{imrt} + \mu^{hd}_{imrt}. \tag{10}$$

Given some vector of slot utilization, the airline-wide team chooses a route wide price component for each flight in each route fo each market they entered. Formally, airline i's profits for route r in market m during period t is given by:

$$\Pi_{imrt}(p_{imrt}, \mathbf{p_{-imt}}, \mu_{\mathbf{imrt}}, \mu_{\mathbf{-imrt}}, \mathbf{f_{imt}}, \mathbf{f_{-imt}}) = \tag{11}$$
$$(\sum_{hd} [\left(p_{imrt} + \mu^{hd}_{imrt}\right) \cdot share^{hd}_{imrt}(p_{imrt}, \mathbf{p_{-imt}}, \mu_{\mathbf{imrt}}, \mu_{\mathbf{-imrt}}, \mathbf{f_{imt}}, \mathbf{f_{-imt}}) \cdot MS_m - MC^f_{imrt} \cdot f^{hd}_{imrt}],$$

where p_{imrt} is airline i's route-level price component; μ^{hd}_{imrt} is airline i's hour/day hd specific price component; $share^{hd}_{imrt}$ is airline i's hd market share; $\mu_{\mathbf{imrt}}$ is the vector of all airline i's hour/day specific price components; $\mu_{\mathbf{-imt}}$ is the vector of all hour/day specific components that all other airlines offer; $\mathbf{f_{imrt}}$ is airline i's daily flight frequencies vector for all hours and weekdays; $\mathbf{p_{-imrt}}$ is a vector with all route-level prices that airline i charges for all other routes in market m, and of all average prices that all other airlines charge for all their routes in market m; $\mathbf{f_{-imt}}$ is a vector with all daily flight frequencies for all other flights that airline i offers in all other routes of market m, and of all flight frequencies that all other airlines offer in all routes of market m; and MC^f_{imrt} is airline i's marginal for one daily frequency. I assume that these marginal costs are the same for all departure hours and weekdays for a given airline-route-market-period tuple.

I remain agnostic to the day-to-day operations team objective function. Intuitively, they may be choosing the hour-day specific component μ^{hd}_{imrt} taking into account operationally specific constraints, such as the available aircraft or crew.

I impose two assumptions on the day-to-day operations team problem. First, Assumption 2 establishes that the airline-wide team behaves as though as the revenue due to the day-to-day operations team spread is equal 0. Intuitively, the airline-wide team may set a target for the average price for a route, and they consider the day-to-day changes as second order. Assumption 3 states that the day-to-day operations team doesn't take into account the price and frequency choices when choosing μ^{hd}_{imrt}. I need this assumption because the data is available at the route level. Intuitively, the day-to-day operations team doesn't care about strategic interactions between their decisions and the airline's decision, as this is the airline-wide team's responsibility.

Assumption 2. *Airline i behaves as though as:*

$$\sum_{hd} \mu^{hd}_{imrt} \cdot share^{hd}_{imrt}(\cdot) = 0$$

Assumption 3. *Let μ^{*hd}_{imrt} be the optimal choice for airline i's day-to-day operations team. For all airlines j, all rmt and all $hd, h'd'$ we have that:*

$$\frac{\partial \mu^{*hd}_{imrt}}{\partial p_{jmrt}} = 0$$
$$\frac{\partial \mu^{*hd}_{imrt}}{\partial f^{h'd'}_{jmrt}} = 0$$

Notice that Assumption 2 implies that we can rewrite airline i's profits from route mrt in Equation (10) can be rewritten as:

$$\max_{p_{imrt}} \Pi_{imrt}(\mathbf{p_{-imt}}, \mu_{\mathbf{imrt}}, \mu_{\mathbf{-imrt}}, \mathbf{f_{imt}}, \mathbf{f_{-imt}}) \equiv$$
$$p_{imrt} \cdot share^{hd}_{imrt}(\cdot) \cdot MS_m - MC^f_{imrt} \cdot \sum_{hd} f^{hd}_{imrt}.$$

Given Equation (10) and the T1EV assumption we can write a closed form rote-airline demand as shown in Proposition 1.

Proposition 1. *Under Equation (10) and the T1EV assumption, the route-airline demand is:*

$$\ln\left(\frac{share_{imrt}}{share_{0mt}}\right) = \overbrace{\beta_p \cdot \frac{1}{\varphi(dist_m)} \cdot_{imrt}}^{\text{price}} + \overbrace{X_{imrt} \cdot \beta_x}^{\text{Other obs charac.}} +$$
$$\ln\left(\sum_{hd} \underbrace{\exp\left[\beta_p \cdot \frac{1}{\varphi(dist_m)} \cdot \mu^{hd}_{imrt} + \beta^{hd}_{irmt}\right] \cdot \#Flights^{hd}_{imrt}}_{\substack{\text{Available flights} \\ \text{by hour of day}}}\right) + \xi_{imrt} \tag{12}$$

Proof. The T1EV errors together with Assumption 3 mean that:

$$\frac{share^{hd}_{imrt}}{share_{0mt}} = \exp\left(X_{imrt} \cdot \beta_x + \xi_{imrt} + \beta_p \cdot \frac{1}{\varphi(dist_m)} \cdot_{imrt}\right) \cdot$$
$$\exp\left(\beta_p \cdot \frac{1}{\varphi(dist_m)} \cdot \mu^{hd}_{imrt} + \beta^{hd}_{imrt}\right) \cdot \#Flights^{hd}_{imrt}.$$

Aggregating up to the route-airline level we get:

$$\frac{share_{himrt}}{share_{0mt}} = \exp\left(X_{imrt} \cdot \beta_x + \xi_{imrt} + \beta_p \cdot \frac{1}{\varphi(dist_m)} \cdot_{imrt}\right) \cdot$$
$$\sum_{hd}\left(\exp\left(\beta_p \cdot \frac{1}{\varphi(dist_m)} \cdot \mu^{hd}_{imrt} + \beta^{hd}_{imrt}\right) \cdot \#Flights^{hd}_{imrt}\right),$$

and taking logs on both sides gets us to the final expression in Proposition 1. □

Next, we can use the airline frequency choice problem to further simplify the demand expression in Equation (12). To do this, I first impose the following normalization assumption on $\beta^{imrt}_{h,d}$ and $\mu^{imrt}_{h,d}$.

Assumption 4. *For all imrt there exists h, d such that $\beta^{h,d}_{imrt} = 0$ and $\mu^{h,d}_{imrt} = 0$. Let $\underline{h}, \underline{d}$ be this h, d.*

This is a relatively mild assumption. In words, this assumption states that there exists one

day/time combination—like for example Tuesdays at 3 AM—that is equally unappealing to consumers regardless of airline, route, market or year.

As in the main text, if we restrict attention to only those routes where airline i has no slots at 80% or 100% usage, then we know that Equation (7) has to hold. This observation and the assumption that marginal costs are the same for all flights with the same route-airline-period tuple lead to the following linear expression.

Proposition 2. *Proposition 1 along with Assumption 4 imply that if (7) holds for all h, d, then:*

$$\ln\left(\frac{share_{imrt}}{share_{0mt}\cdot\sum_{h,d}\#Flights_{imrt}^{h,d}}\right)=\beta_p\cdot\frac{1}{\varphi(dist_m)}\cdot(p_{imrt}+\mu_{imrt}^{\underline{h,d}})+X_{imrt}\cdot\beta_x+\xi_{imrt}$$

Proof. To prove this, we first need to prove the following lemma:

Lemma 1. *Let (i,r,m,t) be a tuple such that Equation (7) holds for all h,d. Then we must have that in equilibrium:*

$$\mu_{imrt}^{hd}=\mu_{imrt}^{\underline{hd}}-\frac{\varphi(dist_m)}{\beta_p}\cdot\beta_{imrt}^{hd}$$

Proof of Lemma 1. Let I_m be the set of all active airlines in marekt m Since (7) holds for all hd, we know that airline i is solving the static unbounded problem for route r in market m during period t. This problem is:

$$\max_{\mathbf{f}_{imrt}}\Pi_{imrt}(p_{imrt},\mathbf{p_{-imt}},\mu_{\mathbf{imrt}},\mu_{\mathbf{-imrt}},\mathbf{f_{imt}},\mathbf{f_{-imt}})$$

As p_{imrt}, $\mathbf{p_{-imt}}$, $\mu_{\mathbf{imrt}}$ and $\mu_{\mathbf{-imrt}}$ are chosen after frequencies, the FOC for f_{imrt}^{hd} is:

$$\frac{\partial\Pi_{imrt}}{\partial f_{imrt}^{hd}}+\frac{\partial\Pi_{imrt}}{\partial p_{imrt}}\cdot\frac{\partial p_{imrt}}{\partial f_{imrt}^{hd}}+\frac{\partial\Pi_{imrt}}{\partial\mu_{imrt}}\cdot\frac{\partial\mu_{imrt}}{\partial f_{imrt}^{hd}}+\sum_{\substack{j\in I_m\\ j\neq i}}\left(\frac{\partial\Pi_{imrt}}{\partial p_{jmrt}}\cdot\frac{\partial p_{jmrt}}{\partial f_{imrt}^{hd}}+\frac{\partial\Pi_{imrt}}{\partial\mu_{jmrt}}\cdot\frac{\partial\mu_{jmrt}}{\partial f_{imrt}^{hd}}\right)=0. \tag{13}$$

Notice that optimality of p_{imrt} implies that the second term is 0. On the other hand, Assumption 3 implies that the third and fifth term are 0. Thus, we have that:

$$\frac{\partial\Pi_{imrt}}{\partial f_{imrt}^{hd}}+\sum_{\substack{j\in I_m\\ j\neq i}}\left(\frac{\partial\Pi_{imrt}}{\partial p_{jmrt}}\cdot\frac{\partial p_{jmrt}}{\partial f_{imrt}^{hd}}\right)=0.$$

Next, we can explicitly compute these derivatives:

$$(p_{imrt}-MC_{imrt}^{q})\cdot\left(\frac{\partial share_{imrt}}{\partial f_{imrt}^{hd}}+\sum_{\substack{j\in I_m\\ j\neq i}}\frac{\partial share_{imrt}}{\partial p_{jmrt}}\cdot\frac{\partial p_{jmrt}}{\partial f_{imrt}^{hd}}\right)\cdot MS_m=MC_{imrt}^{f}. \tag{14}$$

Now let:

$$A^{hd}_{imrt} \equiv \exp\left(X_{imrt}\cdot\beta_x + \xi_{imrt} + \beta_p \cdot \frac{1}{\varphi(dist_m)}\cdot_{imrt}\right)\cdot\exp\left(\beta_p\cdot\frac{1}{\varphi(dist_m)}\cdot\mu^{hd}_{imrt} + \beta^{hd}_{imrt}\right).$$

The T1EV assumption together with Equation (12) mean that:

$$\frac{\partial share_{imrt}}{\partial f^{hd}_{imrt}} = A^{hd}_{imrt}\cdot share_{imrt}\cdot(1-share_{imrt}). \tag{15}$$

On the other hand, to find $\frac{\partial p_{jmrt}}{\partial f^{hd}_{imrt}}$ we need to solve firm j's pricing problem:

$$\max_{p_{jmrt}} \Pi_{jmrt}(p_{jmrt}, \mathbf{p_{-jmt}}, \mu_{\mathbf{jmrt}}, \mu_{\mathbf{-jmrt}}, \mathbf{f_{jmt}}, \mathbf{f_{-jmt}})$$

By Assumption 3, and similarly to Equation 13, we can ignore the indirect effects through $\mu_{\mathbf{jmrt}}$ and $\mu_{\mathbf{-jmrt}}$. Thus, the FOC is simply the derivative with respect to p_{jmrt}, which given the T1EV assumption and Equation (12) is:

$$\begin{aligned} share_{jmrt}\cdot MS_m + (p_{imrt} - MC^q_{imrt})\cdot\beta_p\cdot\frac{1}{\varphi(dist_m)}\cdot share_{jmrt}\cdot(1-share_{jmrt})\cdot MS_m = 0 \\ G(p_{jmrt}, f^{hd}_{imrt}) \equiv 1 + (p_{imrt} - MC^q_{imrt})\cdot\beta_p\cdot\frac{1}{\varphi(dist_m)}\cdot(1-share_{jmrt}) = 0. \end{aligned} \tag{16}$$

Using the implicit function theorem, we know that:

$$\frac{\partial p_{jmrt}}{\partial f^{hd}_{imrt}} = -\frac{\frac{\partial G}{\partial f^{hd}_{imrt}}}{\frac{\partial G}{\partial p_{jmrt}}}. \tag{17}$$

And from (16), the T1EV assumption and (12) we know that:

$$\frac{\partial G}{\partial f^{hd}_{imrt}} = -(p_{imrt} - MC^q_{imrt})\cdot\beta_p\cdot\frac{1}{\varphi(dist_m)}\cdot\frac{\partial share_{jmrt}}{\partial f^{hd}_{imrt}}, \tag{18}$$

where:

$$\frac{\partial share_{jmrt}}{\partial f^{hd}_{imrt}} = -A^{hd}_{imrt}\cdot share_{jmrt}\cdot\frac{1}{1+\sum_{k\in I_m}\sum_{hd}A^{hd}_{imrt}\cdot f^{hd}_{imrt}}. \tag{19}$$

Replacing (15), (17), (18) and (19) in (14) we get:

$$A^{hd}_{imrt} \cdot (p_{imrt} - MC^{q}_{imrt}) \cdot \Bigg(share_{imrt} \cdot (1 - share_{imrt}) + \sum_{\substack{j \in I_m \\ j \neq i}} \frac{\partial share_{imrt}}{\partial p_{jmrt}} \cdot \frac{(p_{imrt} - MC^{q}_{imrt}) \cdot \beta_p \cdot \frac{1}{\varphi(dist_m)} \cdot share_{jmrt} \cdot \frac{1}{1+\sum_{k \in I_m} \sum_{hd} A^{hd}_{imrt} \cdot f^{hd}_{imrt}}}{\frac{\partial G}{\partial p_{jmrt}}} \Bigg) \cdot MS_m = MC^{f}_{imrt}.$$

Since MC^{f}_{imrt} is the same for all hd, it follows that for all $hd, h'd'$:

$$A^{hd}_{imrt} = A^{h'd'}_{imrt},$$

And thus:

$$\beta_p \cdot \frac{1}{\varphi(dist_m)} \cdot \mu^{hd}_{imrt} + \beta^{hd}_{imrt} = \beta_p \cdot \frac{1}{\varphi(dist_m)} \cdot \mu^{h'd'}_{imrt} + \beta^{h'd'}_{imrt}. \tag{20}$$

Combining Assumption 4 with equation (20) yields the desired result. □

From Lemma 1 and Equation (12) we have that:

$$\ln\left(\frac{share_{imrt}}{share_{0mt}}\right) = \beta_p \cdot \frac{1}{\varphi(dist_m)} \cdot_{imrt} + X_{imrt} \cdot \beta_x + \ln\left(\sum_{h,d} \exp\left[\beta_p \cdot \frac{1}{\varphi(dist_m)} \cdot \left(\mu^{hd}_{imrt} - \frac{\varphi(dist_m)}{\beta_p} \cdot \beta^{hd}_{imrt}\right) + \beta^{hd}_{irmt}\right] \cdot \#Flights^{hd}_{imrt}\right) + \xi_{imrt}.$$

Basic algebra leads to the desired result. □

Appendix B Mitigating Potential Hour/Day Bias

In this paper I assume that all flights from a given airline in a given route have the same average utility for consumers. This assumption greatly reduces the airlines' choice set, which is important given that I need to compute the equilibrium of my game for tens of thousands of slot allocations when running my genetic algorithm in Section 8.4. In this appendix, I show evidence that the consumer's preferences for specific departure hours and days load onto the route-airline fixed effects. I also argue that this fact leads to a bias in my demand estimations, and I construct modified route-airline fixed effects that mitigate this bias.

The fact that I don't explicitly model departure hour and day preferences means that they necessarily are part of the route-airline characteristics that the fixed effects captures. Figure 15 presents evidence of this. Here, I show the average willingness to pay that the route-airline fixed effect implies for each flight in my ADRR dataset. To construct this, I gave each flight its route-airline fixed effect based on the flight's route and airline. We can see clear evidence of a morning and evening rush during the weekdays, as well as a drop in value for weekend flights and late night flights. These two facts square well with the argument that the route-airline fixed effects include the departure hour and day preferences.

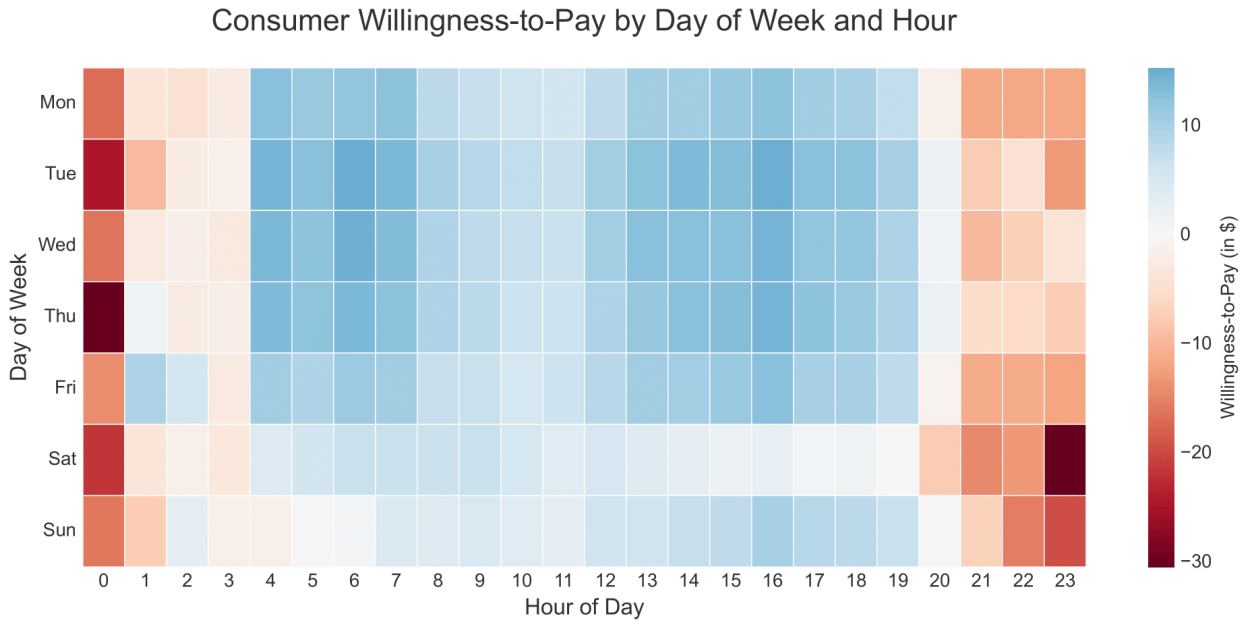


Figure 15: Consumer Willingness to pay for departure hours and days. I compute this by dividing the route-airline fixed effect for each flight by the price coefficient. This plot shows a weekday morning rush centered around 6 AM departures, as well as a weekday evening rush centered around 4 PM departures. In addition, weekend flights as well as late night flights are less desirable.

The loading of departure preferences onto route-airline fixed effects presents a problem for simulation. This is because these fixed effects do not vary with additional flights. Thus, my estimates essentially assume that any flight that the airline adds to a route has the same departure hour and

day as any other flight in that route. This makes some route artificially desirable, and some others artificially undesirable. While this problem can definitely bias my results, the sign of the bias is unclear. On the one hand incentivizes airlines to concentrate their flights on fewer routes, making their profits artificially high. On the other hand, it also shows consumers in those routes as receiving higher levels of welfare than they actually do.

To address this problem, in this appendix I perform and adjustment that aims to correct for this bias. Specifically, I first run the following flight level regression:

$$FE_{imrhd} = \underbrace{FE_i}_{\text{airline component}} + \underbrace{FE_a}_{\text{origin component}} + \underbrace{FE_{a'}}_{\text{destination component}} + \underbrace{FE_h}_{\text{hour component}} + \underbrace{FE_d}_{\text{day component}}.$$

This allows me to extract the departure hour and departure day component from the route-airline fixed effects. I then construct an average departure day fixed effect ($\overline{FE}_h$) and departure hour fixed effect ($\overline{FE}_d$). Finally, I construct a bias-robust fixed effect by replacing the flight-specific hour and day fixed effect with these averages:

$$\widehat{FE}_{imr} = \underbrace{FE_i}_{\text{airline component}} + \underbrace{FE_a}_{\text{origin component}} + \underbrace{FE_{a'}}_{\text{destination component}} + \underbrace{\overline{FE}_h}_{\text{avg hour component}} + \underbrace{\overline{FE}_d}_{\text{avg day component}}. \tag{21}$$

The new fixed effects $\widehat{FE}_{imr}$ all contain the same departure hour and day component, an average of all components. Thus, any differences between these fixed effects come from other route- and airline-specific characteristics.

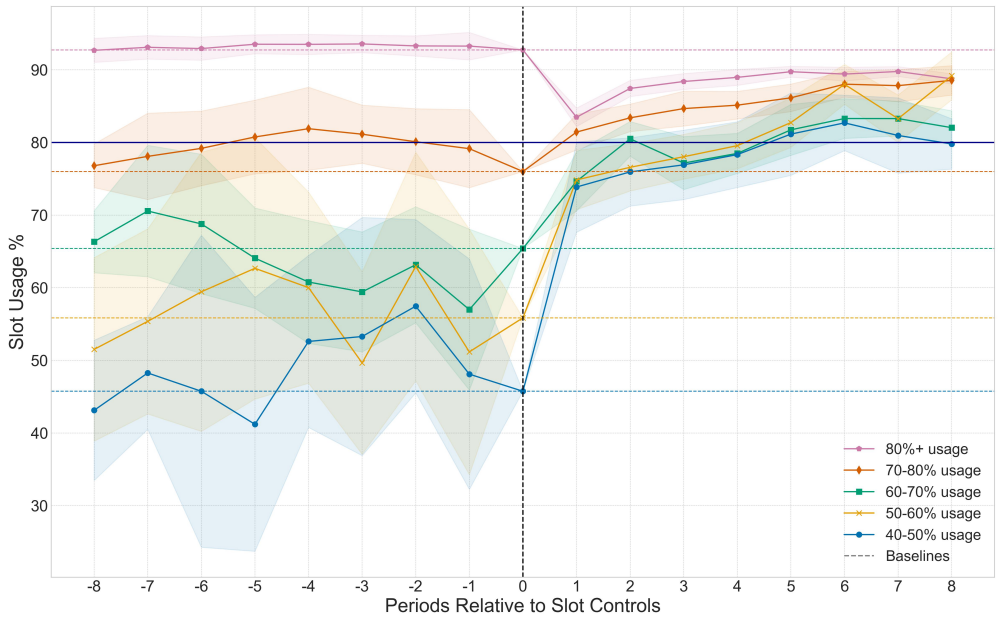
I rerun the baseline specification and the counterfactual of Section 8.3 with these new route-airline fixed effects for the demand. Table 10 summarizes the results. I find that consumer welfare drops significantly more under the counterfactual with these new fixed effects. In addition, airline profits do not grow as much. These two facts combined mean that the loss of the pro-competitive effects of the 80% rule are amplified. This is evidence that the potential bias caused by not explicitly modeling departure preferences actually dampens the differences between the baseline and the counterfactual simulations.

	Baseline	Counterfactual I	Change
Airline Profit	\$16.1 Bn	\$16.4 Bn	2.01%
Consumer Surplus	\$13.8 Bn	\$11.6 Bn	-16.13%
Total Welfare	\$29.9 Bn	\$28.0 Bn	-6.35%
Passengers	490.3 Mn	469.7 Mn	-4.21%
Average Markup	40.05%	40.87%	2.07%
Daily Flights	16,944	15825	-6.60%

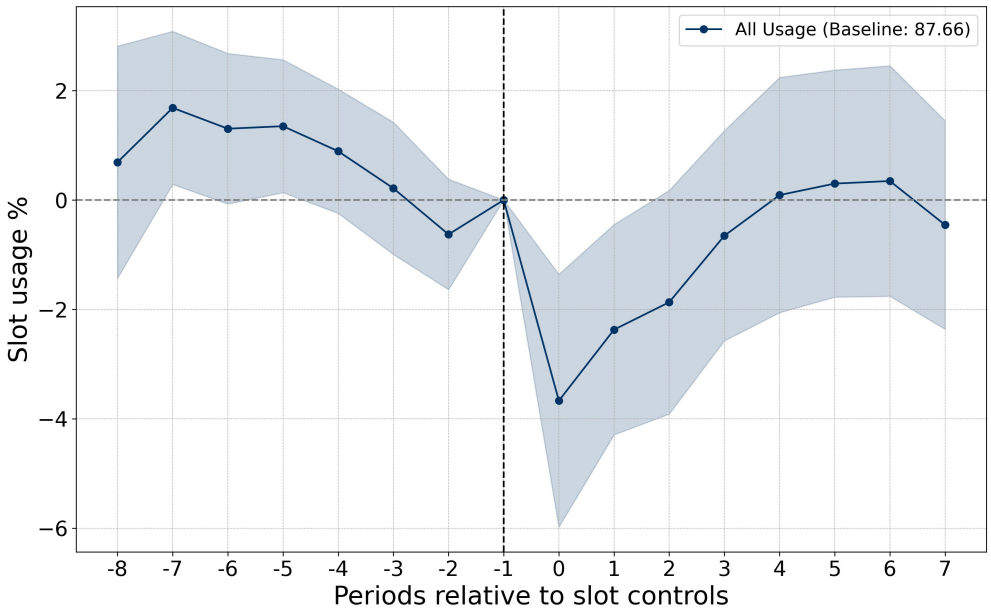
Table 10: Comparison between the baseline simulation in Section 8.1—where airlines are required to fly at least 80% at each slot controlled airport—and the first counterfactual simulation—where the 80% usage requirement doesn't apply. In both cases, the slot allocation is the actual slot allocation of 2016. I adjust the route-airline demand fixed effects according to Equation (21).

Appendix C Event Study Robustness Checks

In this appendix I rerun the event studies from Section 3 while implementing the Borusyak et al. (2024) imputation estimator that controls for potential biases introduced from differential treatment over time. My results remain broadly unchanged.



(a) Time window disaggregation by usage.



(b) Aggregate effect on usage across slots.

Figure 16: Event studies in Figure 2 while applying the imputation estimator in Borusyak et al. (2024).



Figure 17: Event studies in Figure 4 while applying the imputation estimator in Borusyak et al. (2024).

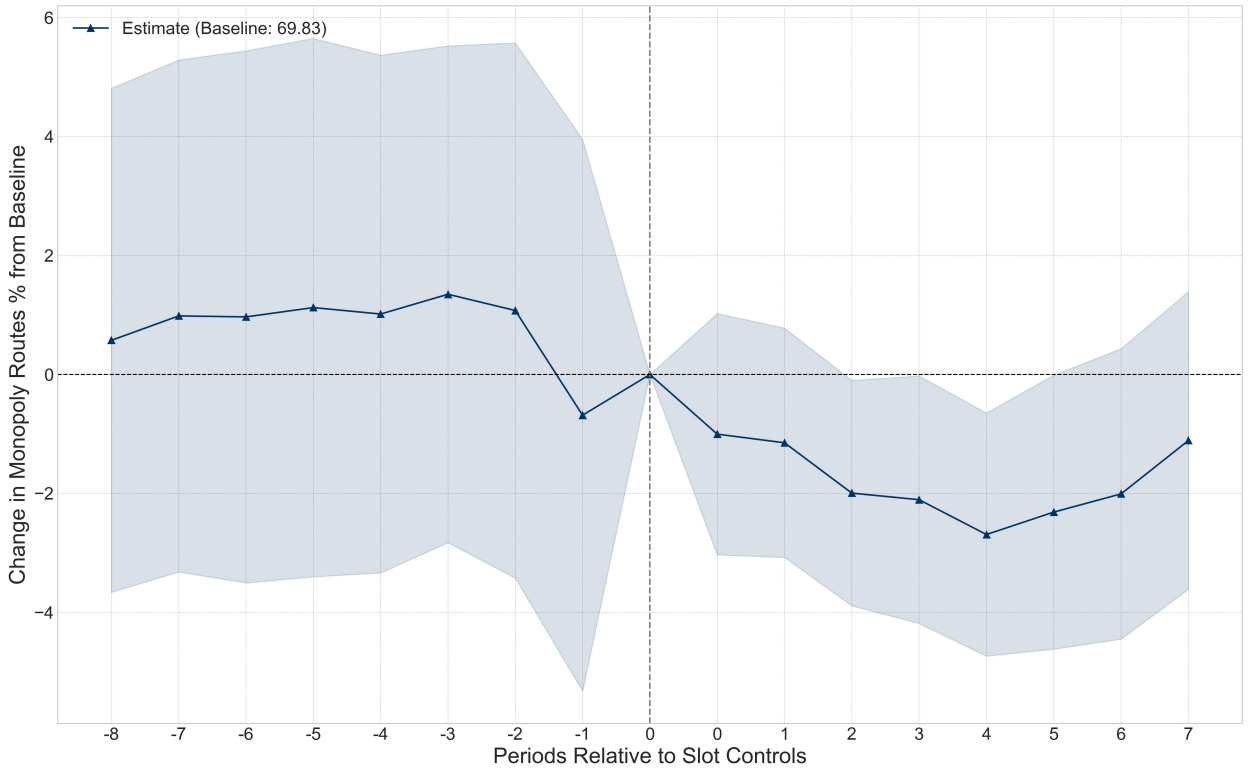


Figure 18: Event studies in Figure 3 while applying the imputation estimator in Borusyak et al. (2024).

Appendix D Slot Matching Across Airports

A pair of slots at two airports a, a' is only useful to airline i for flying route $r \equiv a \to a'$ if the slots are such that if the flight time between airports a and a' is h hours, then the landing slot at a' is scheduled h hours ahead of the departure at a. In this appendix, I first present some strategies airlines have to maximize the chances to get their slot schedules to match up, using examples from my slot database. While these do not completely negate the issue, they do make it less salient. I then run a permutation exercise where I move an arrival slot by one hour and check whether the airline can move flights to/from level 1/2 airports to accommodate the change in the arrival slot time. I find that over half of all these permutations can be successfully performed.

One strategy airlines can employ when they get a new slot is to cancel a flight to an airport that is not slot controlled. For example, in the 2017 Summer season, TAP Portugal got a new Monday 6am departure slot at Frankfurt airport. It used it to fly to Lisbon. As such, the airline needed a 9am arrival slot at Lisbon. It had such a slot, and it was using it to fly to Recife, Brasil. It moved that flight to arrive at 10:30am, and canceled the flight that was arriving to Lisbon from Panama at 10:30am. Since Recife and Panama are not Level 3 airports, this was not an issue.

Full-Service Carriers frequently can count on their fellow alliance members to swap slots, even if they don't have a matching slot themselves. For example, in the 2017 Summer season, IAG got a pair of matching Monday slots at London-Heathrow and Manchester airports. They got the Manchester 8:45am arrival slot from TUI Airways, and the 8:05am departure slot from American. Given that American and IAG are both OneWorld members, it may be possible that there was an agreement between the two airlines after IAG got the Manchester slot from TUI.

For long-haul flights, airlines can adjust their flight speed to conform with slot adjustments. For example, in both the 2016 and 2017 Summer seasons, Air Canada used their Saturday 12pm Heathrow departure slot to fly to Toronto-Pearson. In 2016, this flight used the 7.30pm arrival slot at Pearson. In 2017, despite still departing 12pm from Heathrow, they used the 7pm arrival slot.

Finally, there is always the option to cancel a flight from a slot controlled airport, even if it means losing a slot there. For example, in the 2017 Summer season, Air Berlin got a Wednesday 6:45am Zurich departure slot, which they used to fly to Mallorca. Thus, they needed a 8:45am arrival slot to Mallorca. They had such a slot, which they were currently using for a flight arriving from Munich. Air Berlin canceled that Munich flight and gave up the Munich departure slot, which was taken by Lufthansa to fly to Oslo.

In addition, I also construct a permutation exercise to see how easy it is to adapt to situations where an airline cannot get the slot they wanted. Specifically, I take a the weekly schedule for all European airline groups for 2016, and then for each of them I randomly move a landing slot by between 0 and 1 hour, with a +/- 15 minute leeway. This leeway responds to the fact that coordinators frequently allow for a buffer around the allocated time, as discussed in Section 2. I then try to find a compatible flight to a non-slot controlled airport, so as to swap with the moved flight. I repeated this exercise 1000 times for each airline group. I summarize the results in Table 11. Most airlines can adapt to this change in schedule over 50% of the time.

Airline	**Value**
IAG	59.8%
Lufthansa	64.6%
EasyJet	54.1%
Ryanair	52.8%
AirFrance-KLM	60.2%
Norwegian	49.7%
SAS	62.9%
Air Berlin	52.4%
Alitalia	57.0%
Wizz Air	44.0%
TAP Portugal	46.0%
Finnair	60.2%
Aegean Airlines	54.7%
LOT Polish Airlines	47.6%
Air Europa	38.5%
Volotea	33.9%
Czech Airlines	38.7%
Luxair	32.7%
Air Malta	39.7%

Table 11: Results of the Permutation Exercise

Appendix E Slot Series Building

As I do not have slot ownership data, I reconstruct the slot series that each airline owns based on the ADRR dataset. To do this, I follow a similar method to the one that slot coordinators use to determine which flights belong to which slot series. As airlines have to specify the destination for each airport slot, I construct a slot series by finding flights that have the same scheduled departure during the same day of the week, and also have the same destination. As schedules frequently have to be adjusted over the quarter, many slot series have an allocated time during the quarter but for some weeks I observe a very similar flight during a slightly different time. As coordinators have some tolerance for the departure and arrival times, I fill in missing flights with flights with the same airline and destination that are within 1 hour of the main slot series and have less than 3 occurrences over the season. In addition, I remove flights where the actual departure is more than one hour early/late with respect to the scheduled departure. I exemplify this approach in Table 12.

Origin	Destination	Airline	Date	Scheduled	Actual	Match?
Rome (FCO)	Brussels (BRU)	Alitalia	2017-12-07	14:20:00	14:22:00	**True**
Rome (FCO)	Brussels (BRU)	Alitalia	2017-12-14	14:20:00	14:27:00	**True**
Rome (FCO)	Brussels (BRU)	Alitalia	2017-12-21	14:20:00	14:29:00	**True**
Rome (FCO)	Brussels (BRU)	Alitalia	2017-12-28	14:33:00	14:45:00	**True**
Rome (FCO)	Brussels (BRU)	Alitalia	2018-03-01	14:20:00	14:27:00	**True**
Rome (FCO)	Brussels (BRU)	Alitalia	2018-03-08	14:20:00	15:41:00	**False**
Rome (FCO)	Brussels (BRU)	Alitalia	2018-03-15	14:20:00	14:21:00	**True**
Rome (FCO)	Brussels (BRU)	Alitalia	2018-03-22	14:20:00	14:29:00	**True**
Rome (FCO)	Brussels (BRU)	Alitalia	2018-03-29	13:15:00	13:17:00	**False**

Table 12: Flight matching example. These are flights that correspond to a Wednesday at 14:20 Rome (FCO) departure slot during the Winter season of 2017-2018. The slot operator, Alitalia, uses it to fly to Brussels (BRU). I match all uses where the actual and scheduled departures are both within 1 hour of the 14:20 window as belonging to the same slot series. When a flight's actual or scheduled departure has a difference of more than 1 hour with respect to the predominant slot series' time window, I do not count these flights as part of the slot series. This is what happens with the March 8 flight—it departed 1 hour and 21 minutes late—and the March 29 flight—its scheduled departure was 1 hour and 5 minutes early. I eliminate the March 8 flight from the dataset, as it corresponds to a delay on the existing slot series. The March 29 flight has a random chance of staying in the dataset as its own slot series, as its scheduled departure is significantly different from the existing slot series.

This approach still leaves a considerable number of "slots" with only one observation on my sample. These usually correspond to flights that could in principle belong to a slot series, but I decide against including them because they fall outside the tolerance window that many coordinators use. Keeping them as additional slots would distort the actual slot ownerships, and so I decide to keep a random sample of these single observation flights for the demand analysis, such that the total share of slots that would be granted historic rights is 75% (in line with the figure mentioned by Bichler et al., 2023).

Appendix F Robustness Check to Baseline Computation

In this appendix, I repeat the estimation in Section 8.2 but focusing only on those slots that airlines keep in my data. This aims to model the fact that some slots are more desirable than others. I compare the main text baseline to this modified baseline, and I show that the results do not meaningfully change.

Formally, this Appendix's exercise involves replacing Assumption 5 with the following assumption:

Assumption 5. *Assume airport a is slot-controlled. Then, for all $i \in I$ we have that airline i behaves as though as they face the following restriction:*

$$\sum_{a' \neq a} f^{h,d}_{ia \to a' t} \geq \tilde{S}_{iat}$$

where the left hand side is the sum of all flights airline i has at airport a, and the right hand side is the total number of slots that airline i owns in airport a, and where the usage rate observed in the data is above 75%. Furthermore airline i believes all other airlines face the same restriction.

Intuitively, each airline has a subset of airport slots at each airport that it values enough to fly at all costs. All other airport slots are not important enough for the airline to keep them. An implicit assumption here is that either airline i believes it can recover all slots it lets go this period, or it believes that the slots it loses will not meaningfully alter i's profits if they fall into an opponents hands. Finally, notice that this assumption also turns the dynamic problem in Section 8.1 into a static problem. The only difference is that in this appendix the effective usage requirement will be lower than 80% for some airports, as it will depend on the slot share that the airline lets go in the data.

	Baseline Section 8.1	**Baseline Appendix F**	**Difference**
Airline Profit	\$12.7 Bn	\$12.9 Bn	1.63%
Consumer Surplus	\$14.6 Bn	\$14.4 Bn	-1.46%
Total Welfare	\$27.3 Bn	\$27.3 Bn	-0.02%
Passengers	504.9 Mn	499.8 Mn	-1.02%
Average Markup	36.05%	35.82%	-0.61%
Daily Flights	16,824	16,620	-1.21%

Table 13: Comparison between the baseline simulation in Section 8.1—where airlines are required to fly at least 80% at each slot controlled airport—and this appendix's baseline simulation—where the 80% usage requirement only applies to slots used more than 75% in the data. In both cases, the slot allocation is the actual slot allocation of 2016.

Table 13 summarizes the comparison between the baseline in Section 8.2 and this appendix's alternative baseline. We can see that in all cases, the differences between the two baselines are relatively small, and do not alter the conclusions in any meaningful way. In addition, there is an

82.73% correlation on the frequency choices in both baselines. While around a 25% of airport slots see usage rates below 75% in my data, the vast majority of these have usage rates above 50%. Thus, requiring an 80% usage rate across all slots is a good enough approximation to actual airline behavior.

Appendix G Equilibrium Algorithm and Initial Frequencies

To start the equilibrium algorithm, I propose a random initial frequency profile, where for all routes, airlines start with a random initial frequency between 0 and 4 daily flights. The algorithm then performs an iterated best response scheme, where airlines solve for the optimal prices given the frequencies, and then re-optimize their frequencies given the chosen prices and all the other airline's frequencies. In this Appendix, I show that the equilibrium result does not depend on the starting initial frequency. To do this, I start the algorithm from 100 different seeds, and I record the equilibrium prices and frequencies for each case.

Figure 19 summarizes the coefficient of variation (CV) for prices and frequencies, using a kernel density estimate for both variables. The CV is computed as the standard error over the mean for both variables for each market. As can be seen, the vast majority of observations have almost no variation in either prices or frequencies. Indeed, the average CV is 0.15 for frequencies, and 0.14 for prices. The difference mostly stems from the random draws in the demand shocks, which slightly change the outcome for each seed.

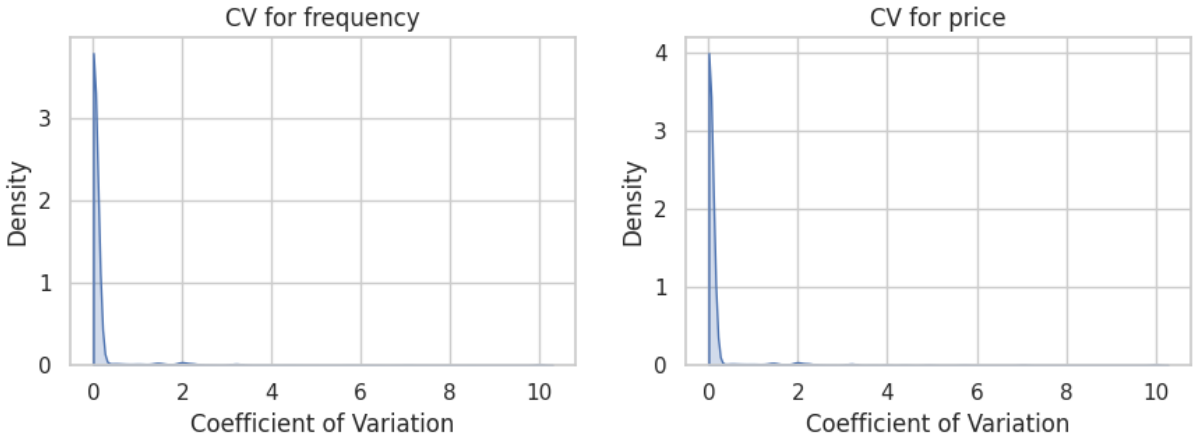


Figure 19: Kernel density estimate for the coefficient of variation for prices and frequencies over 100 random starting points.

Appendix H VCG computation and network effects

Following Krishna (2009), the VCG mechanism in its most general formulation is an efficient mechanism—meaning that the bidder with the highest valuation wins the object—where the payment rule M_i is given by:

$$M_i^{VCG}(\mathbf{x}) = W(\alpha_i, \mathbf{x_{-i}}) - W(\mathbf{x}),$$

where $\mathbf{x}$ is the vector of valuations for all bidders, W is the aggregate welfare function, and α_i is bidder i's lowest possible valuation. In a standard auction context, $\alpha_i = 0$ regardless of the winner's identity, and thus this is equivalent to:

$$M_i^{VCG}(\mathbf{x}) = W_{-i}(0, \mathbf{x_{-i}}) - W_{-i}(\mathbf{x}).$$

This is the payment configuration I use in Section 8.4 to approximate auction payments. Notice however that in my setting it is not true that the excluded airline's valuation is the same regardless of the winner's identity. For example, Lufthansa might prefer that a slot at Frankfurt goes to Air Malta rather than AirFrance-KLM, as the latter has a large intercontinental network which could take passengers away from Lufthansa's direct flights out of Frankfurt. Ignoring these differences could thus introduce a bias into the payments estimation.

Airline Name	**SC Profits %**	**Non-SC Profits %**
Ryanair	83.20%	16.79%
Lufthansa Group	92.43%	7.57%
Air France-KLM	97.12%	2.87%
EasyJet	97.42%	2.58%
IAG (BA-Iberia)	90.96%	9.04%
SAS Group	97.61%	2.39%
Air Berlin	94.72%	5.27%
Norwegian	99.70%	0.30%
Alitalia	100.00%	0.00%
Air Europa	96.42%	3.59%
Wizz Air	30.74%	69.26%
TAP Air Portugal	100.00%	0.00%
Finnair	100.00%	0.00%
Volotea	61.33%	38.66%
Aegean Airlines	96.41%	3.59%
LOT Polish Airlines	100.00%	0.00%
Czech Airlines	100.00%	0.00%
Air Malta	72.41%	27.59%
Luxair	84.52%	15.48%

Table 14: Share of auction profits from routes including Slot Controlled (SC) airports, and not including SC airports by airline.

The fact that I auction all airport slots means this bias is slow. This is because in my 2016 sample, 59 out of the 156 airports in my sample have slot controls in at least one season. As these slot-controlled airports include the most valuable airports, a majority of airlines' profits in my

simulation come from markets where either the origin or destination are slot controlled, as can be seen in Table 14.

Thus, if an airline cannot get slots at slot controlled airports, its capacity to compete is seriously disrupted, which means that the extent to which it cares about which rival gets the slots it will lose is limited. For example, if Lufthansa where to lose all of its slots at Frankfurt, then it will not be able to compete with any airline flying out of Frankfurt, and its profits will always be 0 regardless. This means that while my VCG payments may incur in a bias, that bias is likely to be small for most airlines.

Appendix I Omitted Tables & Figures

Treatment Airports (8)	Gothenburg (W2017), Hamburg (S2016), Hannover (S2019), Krakow (S2019), Paros (S2019), Olbia (S2019), Treviso (S2018), Tromsø (S2019)
Control Airports (44)	A Coruña, Asturias, Bologna, Bratislava, Bremen, Budapest, Chambéry, Cologne, Dresden, Figari, Girona-Costa Brava, Granada, Graz, Ibiza (W only), Innsbruck (S only), Harstad-Narvik, Jerez, Kirkenes, Kiruna, La Palma, Larnaca, Leipzig/Halle, Ljubljana, Luxembourg, Marseille, Menorca (W only), Münster/Osnabrück, Murcia, Nantes, Nuremberg, Pamplona, Ponta Delgada, Pristina, Reus, Rimini, Saarbrücken, San Sebastian, Santander, Sevilla, Sofia, Tenerife-North, Verona

Table 15: List of Treatment and Control airports for the event study. Treatment airports are airports that become slot-constrained (Level 3) between 2015 and 2019, and the table includes the first year-season where they become slot-constrained. Control airports remain Level 2 airports throughout the period. Level 2 airports are airports with manageable levels of congestion, and it is the category slot-constrained airports are before first becoming slot-constrained. For Ibiza, Innsburck and Menorca, I only consider one season, as those airports are Level 3 on the other season.

Airline Group	**# Bases**	**Base List (IATA Codes)**
Ryanair	79	BRU, CRL, SXF, FRA, HAM, CGN, NUE, TXL, BRE, FHH, FMM, NRN, FKB, EIN, BLL, CPH, LUX, RYG, TRF, GDN, KRK, KTW, WMI, POZ, WRO, GOT, NYO, RIX, KUN, VNO, LPA, ACE, TFS, ALC, BCN, GRO, IBZ, MAD, AGP, PMI, SCQ, VLC, VLL, SVQ, BOD, TLS, MRS, BVA, ATH, CFU, RHO, CHQ, SKG, BUD, BDS, BRI, REG, SUF, CTA, PMO, CIY, AHO, CAG, MXP, BGY, TRN, BLQ, TSF, VRN, CIA, FCO, NAP, PSA, PRG, MLA, FAO, OPO, LIS, BTS
Lufthansa Group	30	BRU, FRA, HAM, CGN, DUS, MUC, NUE, STR, TXL, HAJ, LPA, TFS, BCN, IBZ, PMI, NCE, HER, RHO, SKG, CTA, OLB, MXP, FCO, NAP, PRG, SZG, VIE, FAO, GVA, ZRH

Continued on next page

Table 20 – continued from previous page

Airline Group	**# Bases**	**Base List (IATA Codes)**
EasyJet	26	BRU, SXF, HAM, TXL, AMS, CPH, TFS, PMI, BOD, TLS, LYS, NCE, CDG, ORY, LIL, NTE, BSL, CTA, MXP, VCE, FCO, NAP, FAO, OPO, LIS, GVA
IAG (BA-Iberia)	26	BRU, AMS, LPA, ACE, TFS, TFN, ALC, OVD, BIO, BCN, LCG, IBZ, MAD, AGP, MAH, PMI, SCQ, VLC, VGO, SVQ, CDG, ORY, FCO, FLR, LIS, ZRH
Volotea	18	OVD, BIO, PMI, BOD, TLS, MRS, NTE, SXB, ATH, JTR, BDS, PMO, CAG, OLB, GOA, VRN, VCE, NAP
Air France-KLM	16	MUC, AMS, EIN, RTM, BOD, TLS, CFE, AJA, LYS, MRS, NCE, CDG, ORY, LIL, NTE, SXB
Norwegian	15	HEL, CPH, BGO, OSL, TRD, SVG, GOT, ARN, LPA, TFS, ALC, BCN, MAD, AGP, PMI
Air Berlin	13	HAM, CGN, DUS, MUC, NUE, STR, TXL, LPA, TFS, IBZ, PMI, VIE, ZRH
Wizz Air	7	EIN, GDN, KTW, WAW, VNO, BUD, VIE
Air Europa	6	TFN, BIO, BCN, MAD, AGP, PMI
SAS Group	5	CPH, OSL, TRD, GOT, ARN
Aegean Airlines	4	ATH, HER, RHO, SKG
Alitalia	3	MXP, LIN, FCO
airBaltic	2	TLL, RIX
TAP Air Portugal	2	OPO, LIS
Air Malta	1	MLA
Czech Airlines	1	PRG
Finnair	1	HEL
Luxair	1	LUX
LOT Polish Airlines	1	WAW

Table 20: Bases by Airline Group

Airline Group	2015	2016	2017	2018	2019
Lufthansa Group	51,183	55,998	56,368	60,858	61,519
Air France-KLM	31,206	36,703	36,727	31,479	32,340
IAG (BA-Iberia)	24,432	26,471	26,563	25,144	25,903
SAS Group	17,937	19,484	20,472	17,486	17,968
EasyJet	13,670	15,770	15,860	17,978	18,711
Ryanair	13,200	17,281	17,825	19,679	21,756
Air Berlin	13,163	14,846	11,476	0	0
Alitalia	11,936	12,578	12,905	11,030	10,936
Norwegian	11,301	13,170	14,796	13,915	13,497
TAP Air Portugal	7,102	7,897	8,810	8,226	8,033
Aegean Airlines	5,431	5,803	5,760	5,574	6,097
Finnair	4,118	4,731	5,217	4,984	5,199
LOT Polish Airlines	3,187	4,099	5,130	4,966	5,373
Air Europa	2,332	2,570	2,567	2,250	2,318
Czech Airlines	1,443	1,726	1,908	1,621	1,531
Volotea	1,286	1,871	1,881	2,122	2,524
Wizz Air	1,257	1,456	1,655	1,736	2,234
airBaltic	1,103	1,172	1,337	1,266	1,437
Luxair	940	971	1,039	903	923
Air Malta	425	430	490	530	538

Table 16: Airline Schengen Area Slot Holdings by Year (2015-2019). We can see that the top 5 European airlines are among the airlines that own the most slots. In addition, many full-service carriers like SAS, TAP or Alitalia have slot portfolios that greatly exceed their market share. Finally, we can also see how Air Berlin going bankrupt in October of 2017 led to them losing all the slots they had. These were mostly reallocated to Lufthansa, EasyJet and Ryanair.

	(1) Fare	(2) Fare × Dist
Overflight $ per pax	-1.100***	-22.340**
	(0.395)	(8.929)
Fuel $ per pax	0.337***	-0.517
	(0.124)	(2.012))
Overflight $ per pax × Dist	0.0854***	2.198***
	(0.0215)	(0.719)
Fuel $ per pax × Dist	-0.222***	-0.189
	(0.00582)	(0.193)
# Connections at Origin	0.00136**	-0.00550
	(0.000526)	(0.00939)
# Connections at Dest	0.00389***	0.0133
	(0.000550)	(0.0107)
Observations	79,356	79,356
Kleibergen-Paap Rank Test Kleibergen and Paap (2006)	57.96	21.72
Route-Airline & Quarter-Year FE		
Errors Clustered at Market Level		
*** p<0.01, ** p<0.05, * p<0.1		

Table 17: First stage regressions for Table 3.

Star Alliance	SkyTeam	OneWorld
Aegean Airlines	AirFrance-KLM	Air Berlin
Lufthansa	Alitalia	Finnair
LOT Polish Airlines	Czech Airlines	IAG
SAS	Aerolíneas Argentinas	American Airlines
TAP Portugal	Aeroméxico	Cathay Pacific
Air Canada	China Airlines	Japan Airlines
Air China	China Eastern Airlines	LATAM Airlines
Avianca	China Southern Airlines	Malaysia Airlines
Egyptair	Delta	Qantas
Ethiopian Airlines	Garuda Indonesia	Qatar Airways
Air India	Kenya Airways	Royal Jordanian
ANA	Korean Air	SriLankan Airlines
Air New Zealand	Middle East Airlines	
Copa Airlines	Saudia	
South African Airways	Vietnam Airlines	
Asiana	XiamenAir	
EVA Air		
Thai Airways		
United		

Table 18: Major Airline Alliances and Member Carriers, 2015-2019. Schengen-area carriers are highlighted in red. Non-Schengen European carriers (e.g. Turkish Airlines) are not included.

Catchment Area	**# Airports**	**Airport IATA codes**
Ålesund	2	AES, MOL
Belgium	2	BRU, CRL
Berlin-Brandenburg	2	TXL, SXF
Corsica	2	AJA, FSC
French Riviera	2	NCE, TLN
Galicia	3	LCG, SCQ, VGO
Greater Oslo	3	OSL, RYG, TRF
Greater Stockholm	2	ARN, BMA
Île-de-France	3	BVA, CDG, ORY
Lazio	2	CIA, FCO
Lombardy	3	BGY, LIN, MXP
Luxembourg	2	HHN, LUX
Masovia	2	WAW, WMI
Netherlands	3	AMS, EIN, RTM
North Sardinia	2	AHO, OLB
Øresund	3	AGH, CPH, MMX
Greater Palermo	2	PMO, TPS
Rogaland	2	HAU, SVG
Rhine-Ruhr	4	CGN, DTM, DUS, NRN
Tenerife	2	TFN, TFS
Tuscany	2	FLR, PSA
Upper Rhine	4	BSL, FKB, STR, SXB
Upper Silesia	2	KRK, KTW
Veneto	2	TSF, VCE
Vienna-Bratislava	2	BTS, VIE

Table 19: Catchment Areas with more than one airport.

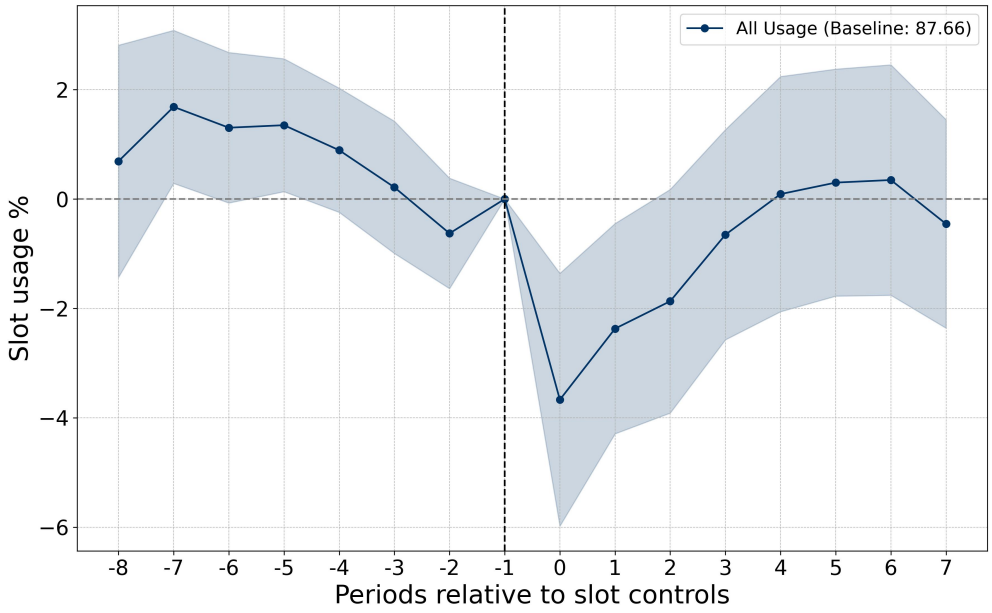


Figure 20: Event study showing the effect of airport slot controls on overall time window usage. My unit of observation is a 5-minute time window during a specific day of the week during a specific season. My unit of time is one week. I cluster errors at the airport level. The value at -1 (the last period before slot controls) is the baseline value. As I am comparing time windows belonging to the same season, there are 36 weeks between the last period without slot controls and the first period with them. The X-axis advances 2 months every 4 weeks due to the composition of the ADRR dataset.

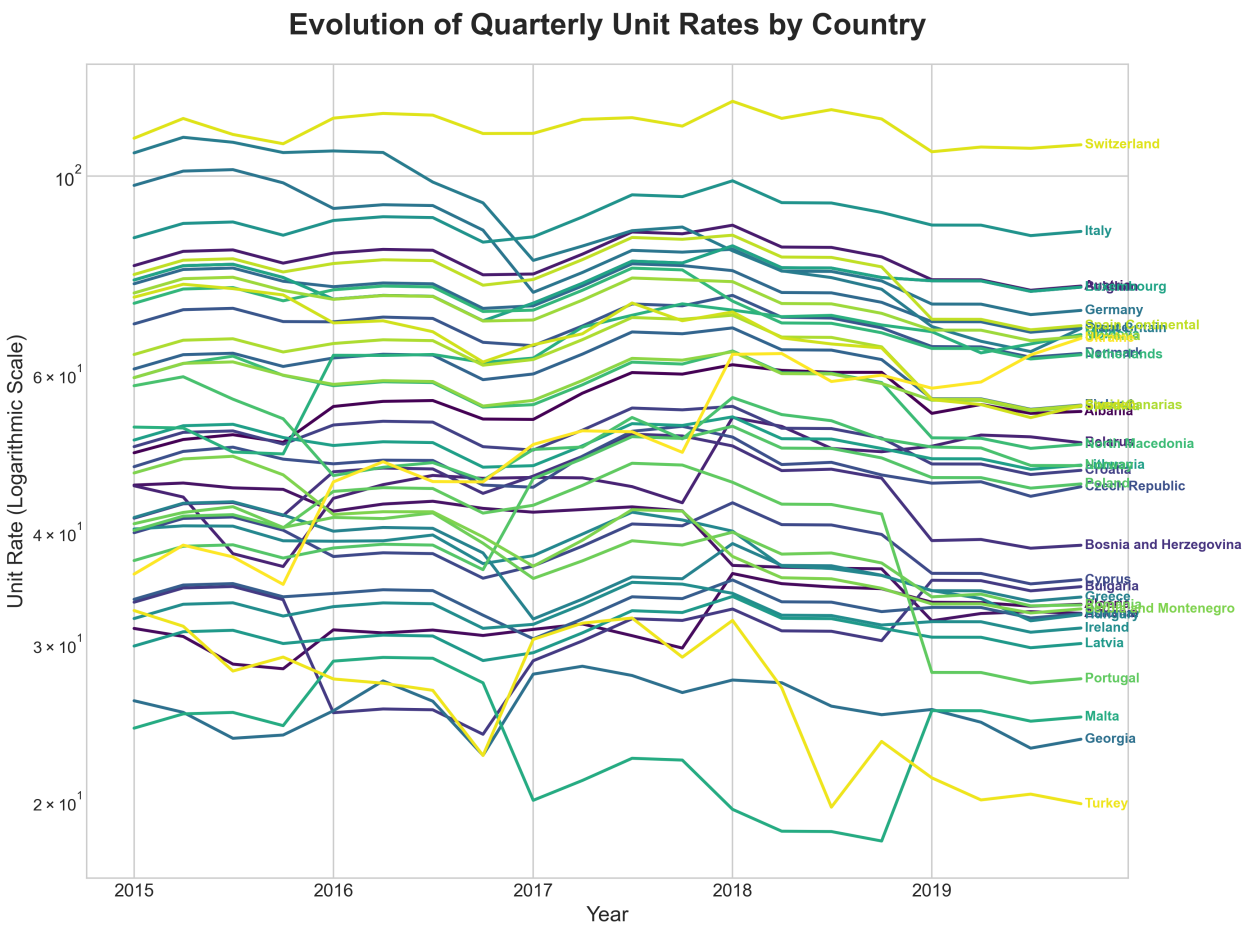


Figure 21: Unit Rate Evolution for Overflight Rights, 2015–2019.

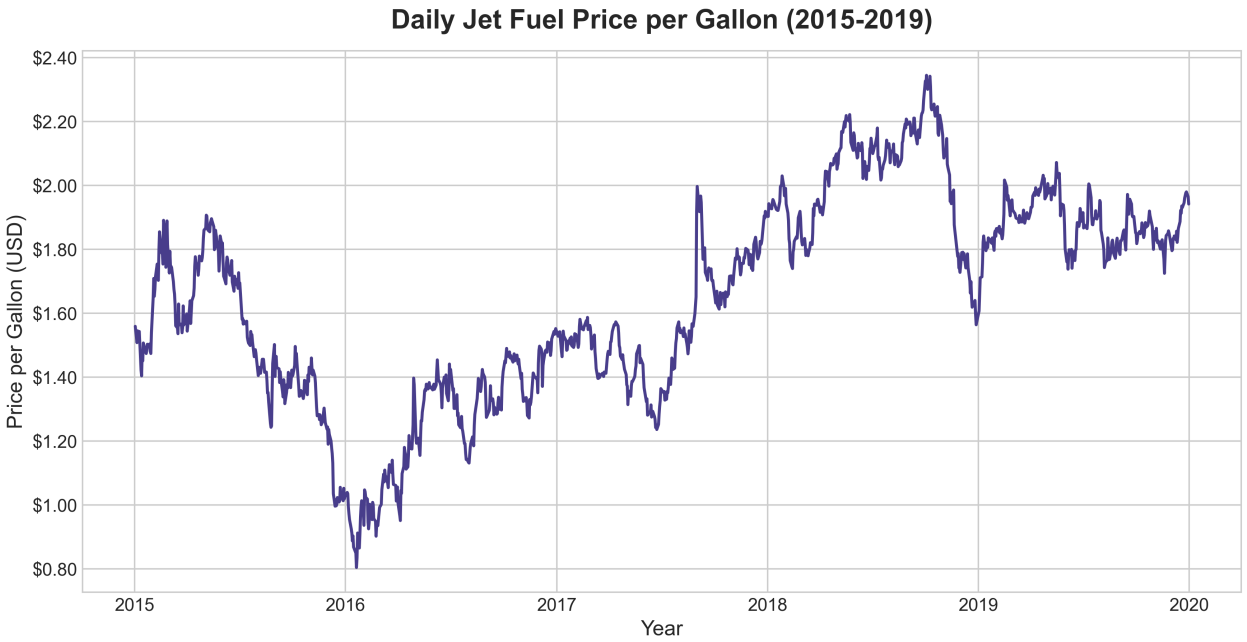


Figure 22: Fuel Price Evolution, 2015–2019.