Essays in Competition Economics

Ratib Mortuza Ali

A dissertation submitted to the Faculty of

the Department of Economics

in partial fulfillment

of the requirements for the degree of Doctor of Philosophy

Boston College
Morrissey College of Arts and Sciences
Graduate School

March 2022

Essays in Competition Economics

Ratib M Ali

Advisors: Professors Julie H Mortimer, Michael D Grubb, Charles Murry

Three self-contained essays explore government regulation in the airline industry, and how such policies affect competition.

The first essay explores the proposed merger between US Airways and American Airlines in 2013, approved by the US Department of Justice (DOJ) under the condition that 104 airport slots ("landing rights") at Ronald Reagan Washington National Airport, DC, be divested to low cost carriers. To investigate the efficacy of the slot divestment, I estimate demand and cost parameters along with bounds on the shadow price of an airline slot, and simulate counterfactual post-merger prices and quantities with and without the regulatory divestment. I find that the merger and associated divestment together increased consumer surplus for markets involving Reagan Airport by roughly 25%. This increase in consumer welfare happened because the median price fell and the quantity of passengers increased. I show that the marginal value of a slot to an airline is decreasing in total slots, validating the DOJ's decision to divest slots from the largest incumbent (US Airways, whose marginal value was $153 per flight) to new

entrants with high valuation (like Southwest, $852). Beyond providing a key input to merger analyses, my approach can also aid in analyzing voluntary exchanges of airline slots, which are subject to DOJ approval due to their perceived anti-competitive effects.

The second essay investigates the impact of airport slots on competition in general. Congestion is managed in high-density airports by capping the number of flights permitted in any given hour and allocating the rights (or slots) to a takeoff or landing among airlines. Airlines must use their slots at least 80% of the time to keep them for the next season. This rule creates a perverse incentive for airlines to hold on to underutilized slots by operating unprofitable flights instead of forfeiting these slots to a rival. Using exogenous removal of slot control at the Newark Airport in 2016, we investigate the lengths at which airlines go to meet the minimum requirements that let them keep the slots while violating what a neutral observer might call the "spirit" of the regulation.

In my third essay, I assess the effectiveness of the gross upward pricing pressure index (GUPPI) in predicting price changes of the 2013 merger between US Airways and American Airlines. I compute GUPPI using only publicly available data, and find that it is close to the observed average increase in price. However, unlike most markets, flights to/from Reagan Airport experience a price drop, likely due to mandated structural remedies; the GUPPI predicts a price increase at Reagan Airport, whereas a full merger simulation correctly predicts a price reduction. I argue that the divergence between GUPPI and, if appropriate, the more accurate predictions of the merger simulation is due to the weaker assumptions made under the simulation. This underscores the fact that while GUPPI, with its restrictive assumptions and low computational burden, can be a good primary screening tool, it does not negate the necessity of employing a more rigorous secondary tool (such as a merger simulation) when assessing mergers.

TABLE OF CONTENTS

Effect of Airport Slots in Competition and Antitrust Policy: Evidence from a Recent Merger

Ratib M Ali

April 12, 2022

Abstract

The Department of Justice (DOJ) approved a proposed merger between US Airways and American Airlines in 2013 under the condition that 104 airport slots ("landing rights") at Ronald Reagan Washington National Airport, DC, be divested to low cost carriers. To investigate the efficacy of the slot divestment, I estimate demand and cost parameters along with bounds on the shadow price of an airline slot, and simulate counterfactual post-merger prices and quantities with and without the regulatory divestment. My estimation assumes that firms maximize profits subject to both flight frequency and price, and that their optimization problem for slot-controlled airports is constrained by their slot endowment. I find that the merger and associated divestment together increased consumer surplus for markets involving Reagan Airport by roughly 25%. This increase in consumer welfare happened because the median price fell and the quantity of passengers increased. I show that the marginal value of a slot to an airline is decreasing in total slots, validating the DOJ's decision to divest slots from the largest incumbent (US Airways, whose marginal value was \$153 per flight) to new entrants with high valuation (like Southwest, \$852). Beyond providing a key input to merger analyses, my approach can also aid in analyzing voluntary exchanges of airline slots, which are subject to DOJ approval due to their perceived anti-competitive effects.

Keywords: merger, structural remedy, competition, antitrust policy, airlines, airport slots.

JEL Classification: D2, D4, K2, L2, L4, L9, R4.

I Introduction

I investigate the competitive effects of the merger between US Airways and American Airlines in 2013, and ask whether (or to what extent) remedies by the Department of Justice (DOJ) mitigated the expected loss in consumer welfare. The DOJ approved a proposed merger between US Airways and American Airlines in 2013 under the condition that 104 airport slots at Ronald Reagan Washington National Airport near Washington, DC, be divested to low-cost carriers (LCCs). Absent this remedy, the newly merged entity would have controlled 591 out of a total of 881 daily slots at Reagan Airport. I evaluate the role of this remedy by calculating the marginal value of an airport slot and comparing the observed divestment to different simulated counterfactual divestment regimes.

Mergers involving major US airlines need prior approval of the DOJ, which can unconditionally approve a proposed merger, categorically deny it, or approve a proposed merger subject to some pre-conditions. Pre-conditions are designed to remedy the negative impacts of the merger on consumers. An important component in any such analysis involving airlines is the allocation to individual airlines of the limited number of gates and slots ("landing rights") at airports. Since 2010, settlements and court orders regarding most proposed mergers involved reallocation or restrictions on gates and slots from one airline to another (this being an example of a structural remedy), but there is little empirical evidence to determine the optimal number of gates/slots to be swapped, or the effects of such divestment on consumer welfare and market structure.[1] My retroactive analysis builds on past work (Kim and Singal, 1993; Peters, 2006; Berry and Jia, 2010; Dobson and Piga, 2013) by incorporating the effects of the structural remedy when evaluating consumer welfare, and varying the structural remedy to investigate alternate outcomes.

I estimate a random utility discrete choice model that flexibly incorporates preferences for characteristics of a flight (following Berry et al. (1995) and Nevo (2001)), including the number of flights on a given route. I assume there are two distinct types of consumers – business and leisure travelers – as opposed to a continuous distribution of heterogeneous consumers (see Kalouptsidi (2012) for how the two approaches are theoretically analogous, and Berry and Jia (2010) for an application of the method). Identification comes from the variation in exogenous variables and instruments used for endogenous variables, choices of which are heavily discussed in related literature (Borenstein, 1989; Berry, 1992; Ciliberto and Tamer, 2009; Berry and Jia, 2010; Ciliberto et al., 2016). This variation allows me to estimate unique coefficients on the product characteristics for a representative business-traveler and leisure-traveler, as well as the proportion of business and leisure travelers in the economy.

I use the demand estimates to compute consumer surplus for consumers flying through Reagan Airport before and after the merger. I find that the merger and the associated divestment together increased consumer surplus by 25.5%, or $7.12 per passenger per one-way flight for travelers in the Reagan market. The merger

[1] Mergers that involve restrictions or reallocation of gates and slots include those between: Virgin America and Alaska Airlines (2016); US Airways and American Airlines (2013); United Airlines and Continental Airlines (2010).

and divestment increased the share of LCCs flying to and from Reagan Airport,[2] the share of direct flights to and from Reagan Airport,[3] and the overall number of flights. Both consumer types prefer LCCs over legacy carriers, direct flights over connecting flights, and more flights to their destination rather than fewer flights, resulting in this increase in consumer welfare. This gain in consumer surplus occurs as median prices fall, number of passengers increase, and the frequency of flights increase. I show that this increase in consumer welfare is due to the divestment.

In order to simulate alternative outcomes, I model the supply side of the operations. One challenge with the model is to treat flight frequency as endogenous; it is common to treat frequency as an exogenous product characteristic (a priori, consumers prefer flights with high frequency since they can find flights closer to their optimal departure/ arrival time). Since the divestiture was a reallocation of slots, it requires firms to adjust frequency, which makes frequency an endogenous variable. Therefore, I assume firms choose both frequency and price to maximize profits.

On the supply side, I estimate cost parameters by assuming firms maximize profits subject to flight frequency and price, allowing me to equate marginal revenue to marginal cost. Ideally, I would assume firms play a three-stage game, choosing network, then frequency (given network choice), and then price (given network and frequency choice). However, while some papers do endogenize airline network choice (Borenstein, 1989; Aguirregabiria and Ho, 2012), to do so and allow firms to choose price for each market is computationally infeasible. Moreover, a separate DOJ settlement (United States v. US Airways Group, 2014) barred the newly merged entity from dropping any routes to small towns, in order to preserve the positive externalities

[2]The divestment only allowed LCCs to apply for the slots.

[3]LCCs operate more direct flights than legacy carriers, owing to their point-to-point business strategy.

that arise from flights to these communities. As a result, I assume the airline network is exogenous, and that firms simultaneously choose price and frequency. Frequency can be thought of as a component of product quality, whose improvement incurs a cost to the firm (Berry and Waldfogel, 2010).

I compute marginal revenue from the estimated demand parameters. Using the firm's first-order condition with respect to frequency, I use variation in the marginal revenue as slot endowments differ across airlines and year to identify bounds on the shadow value of the marginal airport slot, or in other words, the opportunity cost of the marginal flight to an airline. This value is estimated from the supply side, unique for every airline-year pair, ranging between $150 and $550 per flight for legacy carriers, and over $800 for LCCs. The shadow values of the slots confirm that the DOJ decision to restrict the divested slots to LCCs increased market efficiency by allocating the slots to airlines with the highest slot valuations.

To separate the effects of the merger from the divestment, I use the post-merger flight schedule, but pass the ownership of slots from US Airways and American Airlines through to the newly merged entity. My model allows firms to choose the profit-maximizing vector of prices, which then allows me to calculate and compare consumer welfare of the observed merger-with-divestment outcome to a counterfactual merger-without-divestment outcome. A more flexible counterfactual that allows airlines to re-optimize their flight schedule in addition to prices would be more realistic. However, re-optimizing firms' flight schedules involves modeling the entire network of flights for each airline, which is computationally burdensome, and left for future work.

I contribute to the merger literature that investigates the efficacy of structural remedies.[4] Structural remedies are generally thought to be better equipped than conduct remedies[5] at combatting the anticompetitive effects of mergers (Kwoka, 2013), and successful implementation is often possible, as found in the brand divestiture following Johnson & Johnson's acquisition of Pfizer's consumer health division (Tenn and Yun, 2011). However, other works show that such remedies may be futile in shoring up competition (Cabral, 2003), especially if the divestitures are being proposed by the merging entities themselves (Vasconcelos, 2010) to appease DOJ concerns. My unique contribution to this literature is to suggest a framework to evaluate the welfare implications of this and similar mergers, which is possible because slots are relatively identical to one another (unlike brands). In addition, I also compute the marginal value of an airport slot, which informs us of an airline's value of an additional airport slot.

Even though the existing literature is rich in entry and merger analysis, it is nascent when it comes to airport slots. Exceptions include theoretical work on voluntary exchange of slots (Reitzes et al., 2015), which corroborates my findings that reallocating slots from large holders to small ones enhances consumer welfare, and slot allocation as a mechanism design problem (Schummer and Vohra, 2013; Polsby, 2001). My research shows how to incorporate airport slots to the standard demand estimation framework, and how to empirically model firm decisions to reallocate flight frequency in the face of regulatory changes.

My work complements literature on entry in the airline industry (Ciliberto and Tamer, 2009; Ciliberto et al., 2016). Entry decisions can be made along the extensive

[4] Structural remedies alter a firm's structural composition by forcing them to divest assets or brands, and is common for horizontal mergers.

[5] Conduct remedies regulate the conduct of the firm with its competitors, such as an information firewall between brands, or in the case of vertical mergers, forced non-discrimination policies between competing clients of a firm, when one of the clients merges with the firm in question.

margin (airline entering a new route) or the intensive margin (increasing the number of flights on a route). Following the US Airways-American Airlines merger, slots were divested away from the merged entity. However, a separate settlement with the DOJ prevented the merged entity from closing low-density routes to small communities; the data corroborates that no routes were completely shut down by the merged entity. Therefore, route reoptimization by the merged entity comes from decreasing frequency on routes without shutting them down. While existing literature on dynamic entry focuses on analyzing the extensive margin (Ciliberto and Tamer (2009) does so in the context of the repeal of competition restriction for Dallas Love Airport, and Ciliberto et al. (2016) for the same US/American merger), I model airline decisions along the intensive margin to reallocate flights from one pre-existing route to another following an exogenous shock in the number of slots owned by an airline.

The rest of the paper is organized as follows. Section II explains the market structure and the divestment regime in detail. Section III explains the data, and highlights how the data structure informed the modeling choices. Section IV describes the theoretical model, and discusses the choice of instruments and approach to estimation. Section V discusses my findings, and Section VI concludes.

II Background

The merger between US Airways Group and AMR Corporation – the parent companies of US Airways and American Airlines – was announced in February 2013, following the bankruptcy of American Airlines in 2011. Under Chapter 11 protection, American Airlines was seeking a potential merger partner from among its creditors (including US Airways). The merger was approved shortly after the proposal, in March 2013, in the US Bankruptcy Court. The US DOJ, along with six states (Arizona, Florida, Pennsylvania, Tennessee, Texas[6], and Virginia) and the District of Columbia filed a petition in the US District Court in August 2013 seeking to block the merger. Plaintiffs reached a settlement in November 2013, with the merger being approved, subject to some structural remedies – specifically, the divestment of 104 airport slots from the merged entity (who would control 591 of the 881 total daily slots at Reagan Airport absent this remedy) to LCCs.

The new corporation formed in December 2013, and the airline was required to complete the divestment within 180 days. The two brands were given a single operating license in April 2015, and industry insiders noted that the two brands kept their crew and resources separate until the last US Airways flight flew in October 2015.

II.1 Slots

A slot is a permission to perform one departure or one arrival from an airport within a one-hour window on a given day. Slots can be considered as a more granular version of a gate – an airline can use a gate to schedule multiple departures and arrivals on a given day, but only one departure or arrival is allowed per slot. One can loosely say that one gate has many slots associated with it (this is technically incorrect, since a slot is not tied to a specific gate).

Only three airports in the United States currently have slot restrictions – John F Kennedy Airport and LaGuardia Airport in New York City, and Ronald Reagan Washington National Airport (DCA). The existence of slots is an indication that these

[6]Texas dropped out of the petition in September when the merging firms promised to maintain the headquarters of the newly merged entity in Dallas, Texas, and to continue their flights to the smaller towns in Texas.

airports are congested; the slots are allocated by the Federal Aviation Administration (FAA), as opposed to the individual airport authorities, who allocate airport gates. No other US airports have such restrictions on use of airport gates, but they are common at congested international airports (including Toronto Pearson, London Heathrow, Paris Charles de Gaulle, Frankfurt, Beijing Capital, Shanghai Pudong, and Sydney).

The settlement required the following structural remedies from the merged firm (United States v. US Airways Group & AMR Corporation, 2013)[7]:

Table 1: Structural Remedies Required

Airport Name	Type	Slot remedy	Gate remedy
Reagan, DC	Gates & slots	104 slots	Up to 5
LaGuardia, NYC	Gates & slots	34 slots	2
O'Hare, Chicago	Gates	—	2
Los Angeles	Gates	—	2
Love Field, Dallas	Gates	—	2
Logan, Boston	Gates	—	2
Miami	Gates	—	2

The divested slots from the two slot-controlled airports were allocated to JetBlue, Southwest, and Virgin America (see Table 2 for details). American Airlines earned $425 million from the mandatory sale of slots (Maxon, 2014), or about $3 million per slot, which is between the upper bound of the shadow value for the newly merged entity (the seller of the slots) and below the lower bound of the shadow value for the low-cost entrants (the buyers of the slots).

[7]In addition, although not formally part of the DOJ remedy, the US Department of Transportation reached a separate agreement with American Airlines that they will use all of their commuter slots at DCA to serve small, medium, and non-hub airports (i.e. airports that enplane less than 1% of annual domestic passenger enplanement) for at least five years (United States v. US Airways Group, 2014).

Table 2: Allocation of divested slots

Reagan Airport: 104 slots divested	
Southwest	56
JetBlue	40
Virgin America	8
LaGuardia Airport: 34 slots divested	
Southwest	22
Virgin America	12

Note: 881 slots at Reagan Airport per day, or (up to) 60 slots per hour; 71 slots per hour at LaGuardia Airport.

Table 3: Total Slots Allocation at Reagan Airport

Airline	Slots before merger	Change in data	Change by judgment	Slots after merger
American	118	-100	-104	18
Delta	104	—	—	104
JetBlue	20	+40	+40	60
Southwest*	6	+76	+56	82
United	82	—	—	82
US Airways	473	—	—	473
Virgin America	2	+8	+8	10
Others	71	-19	—	52
Total[8]	876	+5	—	881

*Southwest also acquired AirTran during this period, gaining 20 additional slots at DCA

III Data

I obtained data on operational ownership of each slot through requests under the Freedom of Information Act (FOIA) from the Federal Aviation Administration (FAA). The information is at an hourly level, aggregated to a quarterly level because of the restrictions placed by the quarterly structure of the DB1B database. Information on required divestment of slots is taken from the summary judgment for the merger (United States v. US Airways Group & AMR Corporation, 2013). Table 3 provides a breakdown of the slot allocation at Reagan Airport before and after the divestment.

[8] Discrepancies between the data and judgment arise from the difficulty in ascribing regional airlines to a ticketing carrier. For instance, whether an "Air Wisconsin" flight is operated for

The slots data are supplemented by the publicly available DB1B database, which contains a 10% sample of all tickets issued for travel originating and ending in the United States. Each record contains the fare paid, origin, destination, all connections, miles flown, number of passengers ticketed together, and the ticketing carrier, at a quarterly level. I drop all observations where the market fare is less than $50,[9] and those with more than two segments per journey.

Since the focus of my study is on the effects of the merger at Reagan Airport, I consider all itineraries in the contiguous United States originating or terminating at Reagan Airport, keeping one-way tickets and splitting roundtrips as two one-way journeys. While most research on airlines focus on medium-to-large airports only, my work involving slot allocation requires me to consider all routes from Reagan Airport. I exclude routes that use Reagan Airport only as a connection (for example, Boston-Washington-Austin), since a vast majority of passengers traveling through Reagan Airport originate or terminate their journey at Reagan Airport.[10] My sample contains 571,202 passengers across four years. A summary of my observations are listed in Table 4.

I use data from two quarters – the third quarters of 2012 and 2015. Since the merger was proposed in February 2013, I consider the third quarter of 2012 as my pre-merger period. The post-merger period was chosen to be the third quarter of 2015, which captures the period when the slot divestment was fully consummated (by 2014). The third quarter is chosen to control for seasonality in demand for air

American Airlines or United Airlines is manually determined.

[9]This removes award tickets from my sample.

[10]84.01% of all passengers at Reagan Airport either originate or terminate their journeys at Reagan Airport, which is large compared to the same statistic for established hubs: Dallas-Forth Worth (48.12%), Atlanta (35.61%), or Charlotte (32.28%). Incorporating itineraries connecting through Reagan Airport will only add computational burden and consideration for substitution patterns between other connection points, but will not enrich the results in any meaningful way.

Table 4: Summary Statistics

	Year 2012, Q3 (Before merger)		Year 2015, Q3 (After merger)	
Variables	Mean	Std. Dev.	Mean	Std. Dev.
Fare	243.80	127.66	238.33	121.05
Distance (miles)	1,061.26	674.12	1,043.12	677.95
Population (millions)	4.61	2.02	4.82	2.16
Layover	0.48	0.50	0.40	0.49
Legacy	0.85	0.35	0.79	0.41
Tourist	0.20	0.40	0.18	0.38
Carriers in the Market	4.65	1.33	5.07	1.34
Observation	124,522		155,691	
Markets	293		291	

All means are significantly different between 2012 and 2015 at a 99% confidence level.

Population is the geometric mean of the populations of the two end-cities, measured in millions.
Layover is an indicator for whether a passenger has layovers (layover = 1) or flies direct.
Legacy is an indicator for whether a passengers uses a legacy carrier (legacy = 1) or low-cost carrier.
Carriers in the Mkt is the number of carriers that serve a given market.
Markets is the number of unique markets being served.

travel.

I define a market as an unordered city-pair in a year (for example, Washington-Boston-2012), leaving me with 1,288 markets. In line with literature (Berry et al., 2006), the market size is the geometric mean of the MSA population of the end-point cities.

I distinguish products on the basis of the carrier and whether the journey is nonstop or connects through another airport (I make no distinction between the identities of anyn connecting airports). Even at that level of definition, a Boston-

Washington-JetBlue-nonstop product that costs \$120 might be vastly different from the same product costing \$400. The lower-priced ticket might be purchased six months in advance, and taking place on a Tuesday afternoon, as opposed to the higher-priced ticket purchased a week in advance for a Friday evening departure. Since I don't observe features such as day/time of purchase, departure, or arrival, I separate my products into progressive fare bins (following Berry and Jia (2010)) of \$200 to capture the differences in these two products. Therefore, a product in a market is an unordered airport-pair-carrier-layover indicator-fare bin pentuple.

I generate the number of flights on a given route at the market-carrier level using the T100 database, which contains the population of all airplanes flown at a monthly level. I aggregate all the direct flights in the market offered by the carrier. To incorporate the number of connecting flights, I compile all possible connecting itineraries along the route offered by the carrier, and include the minimum of the flight frequencies from the two legs as the number of flights along the route. I add the minimum flight frequencies from all possible connecting routes to the number of direct flights to calculate the total number of flights offered by a carrier in a market.[11]

IV Model

I assume firms simultaneously choose both prices and flight frequencies, and then demand is realized. The firm choice can be thought of as a simplification of a three-step decision process for firms, where firms choose network, frequencies, and prices in

[11]For instance, for Boston-Washington-2012-American quadruple, I consider all possible connecting itineraries. For an itinerary Boston-Charlotte-Washington, I take the minimum of the direct flights on Boston-Charlotte and Charlotte-Washington as the number of possible flights on the Boston-Charlotte-Washington itinerary. I add the minimum frequencies for all possible Boston-x-Washington connections to the number of direct flights between Boston-Washington, to obtain the total number of possible flights offered by the Boston-Washington-2012-American quadruple.

three distinct steps. As discussed earlier, the response to the merger by the merged entity was reducing frequencies along routes, as opposed to shutting down routes entirely. Table 5 shows that the total number of markets or the number of markets served by US Airways or American Airlines did not change meaningfully following the merger.[12] The reason for this is mostly legal – in a separate settlement with the DOJ, the newly merged entity promised not to terminate any routes serving small communities (United States v. US Airways Group, 2014). The relevant focus of my study to investigate the efficacy of slot divestment, therefore, is along the intensive margin, and I assume the network as given. Therefore, I assume firms choose price and flight frequency simultaneously.

Table 5: Number of markets served

	Before merger	After merger
Total Markets	293	291
Markets served by US Airways	156	
Markets served by American Airlines	139	
Markets served by US or AA	210	212

In the second stage, consumers buy tickets and demand is realized. Given my focus on the use of slots and its effects on demand, I focus my analysis on journeys originating or terminating in Reagan Airport (DCA).

IV.1 Demand

Each consumer, i, traveling in market t, either chooses to consume one product j, or not. A market t is a nondirectional-city-pair–year triple; a product j is a nondirectional journey between the pair of cities served by a given airline, and differentiated by whether a journey is direct or connecting. A market contains all

[12]Moreover, the entry by low-cost carriers did not result in service to a new market – entry by LCCs were exclusively in markets already served by incumbents.

the products in the consumer's choice set. For example, in the Boston-Washington-2012 market: Boston-to-Washington and Washington-to-Boston are both the same product; Boston–JFK–Washington and Boston–Washington are different products; Boston–JFK–Washington and Boston–Charlotte–Washington are the same product; Boston–Washington by American and by JetBlue are two different products; and Boston–Washington–Boston roundtrip with the same carrier is considered as one product with two observations. Lastly, passengers pay different prices for the same flight; I capture this variation by assuming that passengers who purchase from different price bins on the same flight consume different products from one another.[13] To achieve this, I separate each product into progressive bins (Berry and Jia, 2010), with an increment of $200. Therefore, an examplar product in the Boston-Washington-2012 market will be: "JetBlue-direct-priced between $50 and $250," while another product might be "American-connecting-priced between $450 and $650."

I use a discrete-choice demand model (Berry & Jia, 2010), where each consumer choooses product j in market t to maximize their utility, u_{ijt}, given by:

$$u_{ijt} = x_{jt}\beta_r - \alpha_r p_{jt} + \beta_r^f f_{jt} + \xi_{jt} + \epsilon_{ijt} \tag{1}$$

where,

- $r \in \{l, b\}$ is consumer i's type (leisure or business traveler)
- x_{jt} is a vector of product characteristics;
- β_r is a vector of "tastes for characteristics" for consumers of type r;
- α_r is the marginal disutility of a price rise for consumers of type r;
- p_{jt} is the product fare;

[13] Consumers may pay different prices based on the day/date/time of the flight, date/time of purchase, or for premium services like priority boarding or extra legroom. I don't observe any of this variation, and therefore take into account this variation by splitting passengers into progressive fare bins.

- β_r^f is the utility from an additional flight available on a route;
- f_{jt} is the number of flights available on the route by a specific carrier, differentiated by direct or connecting;
- ξ_{jt} is the unobserved (to the econometrician) product characteristic; and
- ϵ_{ijt} is the iid logit error.

In this model, instead of each consumer having a unique taste parameter, I bin consumers into $r = 2$ types. Given the stylized fact that consumer tastes for connections, scheduling, and prices are correlated (some consumers are more sensitive to price and less sensitive to the duration/timing of journey and the number of connections, while other consumers exhibit the opposite tastes), a discrete consumer-type model captures these correlations neatly without estimating the full variance-covariance matrix for the continuous random coefficients model.[14] In line with our understanding of consumer heterogeneity in air travel, these consumers can be categorized as leisure or business travelers.

The vector of product characteristics, x_{jt}, include the logarithm of distance between the two airports, a dummy for whether the carrier is legacy or low-cost, distance to the closest alternate airport,[15] a layover indicator, whether the destination is a predominantly tourist destination (measured by a indicator for flights to and from Las Vegas or Florida), and the carrier-level standard deviation of price on a route.

[14] The type-specific parameters and the proportion of business-type passengers are identified by rationalizing the substitution patterns among similar products when the menu of products differs across markets. A continuous random-coefficient logit model estimates k means and $k(k+1)/2$ elements on the variance-covariance matrix (although the full covariance matrix is rarely estimated in reality); a type-specific logit model, à la Berry and Jia (2010), estimates $r \times k$ consumer taste parameters and the $r - 1$ parameters for proportions of each type, which is less than the number of elements measured for the random-coefficient model when $r = 2$, while capturing the consumer correlation in taste.

[15] For the market Raleigh-Washington, each endpoint has an alternate airport closest to it. Since my sample is limited to all markets involving Reagan Airport, I take the distance from the non-Reagan airport (in this example, Raleigh Airport) to its closest airport (Greensboro Airport) as the distance to the closest alternate airport (71 miles).

Price and flight frequency are written out separately to highlight that they are endogenous and need to be instrumented. Flight frequency is a stand-in for the quality of match between a consumer's desired departure/arrival time and the departure/arrival time offered by the product (as in Hotelling (1929)). Frequency can be thought of as a measure of product quality, whose improvement benefits consumers, but also incurs a cost to the firm (Berry and Waldfogel, 2010), calling for it to be treated as an endogenous variable.

A consumer i chooses product j in market t if their utility from j exceeds the utility from any other product j',

$$u_{ijt} > u_{ij't}, \qquad \forall j, j' \in J$$

where $j = 0$ denotes the outside good, and utility from the outside good is normalized to zero ($u_{i0t} = 0$).

This results in the standard logit share equation:

$$s_{jt}(x_t, p_t, f_t, \xi_t, \theta_d) = \sum_r \gamma_r \frac{\exp\left(x_{jt}\beta_r - \alpha_r p_{jt} + \beta_r^f f_{jt} + \xi_{jt}\right)}{\sum_{j=0}^{J} \exp\left(x_{jt}\beta_r - \alpha_r p_{jt} + \beta_r^f f_{jt} + \xi_{jt}\right)} \tag{2}$$

where γ_r is the proportion of type-r consumers, and $\delta_{rjt} = \exp\left(x_{jt}\beta_r - \alpha_r p_{jt} + \beta_r^f f_{jt} + \xi_{jt}\right)$ is the mean utility of a product enjoyed by a representative type-r consumer.

I run a modified version of the standard BLP logit demand estimation model; the modification (Kalouptsidi, 2012) allows me to leverage the fact that I have (two) distinct consumer types, and allows me to recover γ_r, the proportion of type-r consumers.

IV.2 Supply

While estimating demand, I bin each route by a progressive-price bin to capture the heterogeneous nature of the product by price. This distinction is moot for the supply side, since my object of interest is the marginal decision to operate a flight, and I have no way of separately identifying the number of seats on a flight being sold to a specific progressive-price bin. Therefore, the supply side product is a route-year-carrier-connecting status quadruple, subscripted by the letter k.

Profits for the multiproduct carrier c are:

$$\Pi_c(\mathbf{p}, \mathbf{f}) = \sum_{\forall k \in K_c} \Pi_k(\mathbf{p}, \mathbf{f}, MC_k) = \sum_{\forall k \in K_c} \Big(M s_k(\mathbf{p}, \mathbf{f}) \cdot p_k - f_k \cdot MC_k \Big) \tag{3}$$

where, M is the market size, MC_k is the marginal cost of a single *flight*, and K_c is the set of products sold by carrier c, and f_k is the frequency offered by a carrier-connecting status on a route, which includes both direct flights, d_k, on the route as well as connecting flights, $conn_k$. and $f_k = d_k + conn_k$. Boldface indicates the vector of the frequencies and prices of all the products in the market.

Firm c maximizes profits with respect to price and frequency, subject to the slot constraint:

$$\max_{\{p_k, f_k\} \forall k \in K_c} \sum_{\forall k \in K_c} \Pi_k(\mathbf{p}, \mathbf{f}, MC_k) \quad \text{s.t.} \sum_k d_k \leq S_c \tag{4}$$

where S_c is the total number of slots available to firm c. The slot constraint is defined for direct flights, d_k, because the universe of direct flights encompass all the physical flights departing from/arriving to Reagan Airport (DCA), including all connecting flights from it. I set up the Lagrange function for firm c as follows:

$$\mathcal{L}_c = \sum_{\forall k \in J_c} \Pi_k\big(\mathbf{p}, \mathbf{f}, MC_k\big) + \lambda_c(S_c - \sum_k d_k) \tag{5}$$

where λ_c is the shadow value of a slot operated by airline c. The firm has two sets of first-order conditions (FOCs), with respect to price and frequency. I only need to consider the FOC with respect to frequency to estimate the shadow value of a slot.

IV.2.1 FOC with respect to frequency, and the Lagrange multiplier

The FOC with respect to frequency is:

$$\frac{\partial \mathcal{L}_c}{\partial f_k} = \underbrace{M \sum_{\forall i \in J_c} \frac{\partial s_i}{\partial f_k} p_i}_{MR_k} - MC_k - \underbrace{\frac{\partial d_k}{\partial f_k}}_{=1} \lambda_c = 0 \tag{6}$$

where J_c defines hte set of carrier c's products (specific to price bins), denoted by subscript i. Since $d_k = f_k - conn_k$, $\frac{\partial d_k}{\partial f_k} = 1$.

An interesting result of this paper is my ability to identify the shadow value of the marginal slot to an airline, λ_c. I can recover a different λ_k for each flight, and in theory, all λ_k's will be equal for a given carrier. Therefore, it is appropriate that I instead estimate a single λ_c for a given carrier, because the slot endowment is determined and binds at the carrier level.[16] In other words, the multiplier is specific to a carrier because the marginal decision to allocate a slot to a route is determined at the carrier level. For example, American's marginal slot will be allocated to the route with the highest marginal profit, which may (not) be the same route or the same marginal profit if the marginal slot was being allocated by JetBlue. Therefore,

[16] This also has the added advantage of smoothing out noise in the data, and logistical and operational barriers that may cause the shadow values to not be equal across all products for a given carrier.

the shadow cost of the marginal slot is meaningful at the carrier level.

The shadow value of a slot informs us of how an airline considers tradeoffs between competing routes within their choice set. This shadow value also captures the capacity of an airport to support more flights; since λ_c captures an airline's value of the marginal slot, an airport authority can estimate this shadow value to determine whether airport expansion will cause profit-maximizing firms to increase frequency. Essentially, this shadow value is of interest in itself because it captures a firm's incentives to deviate from the observed equilibrium menu of route frequencies.

In the existing literature, this shadow cost of the flight would be implicitly present in the marginal cost of all flights originating from or terminating at a slot-controlled airport (through the constant term). I capture the shadow value of the marginal slot to an airline through fixed effects, separately from the variable cost that is correlated with distance, which is important in modeling the firm's decision and subsequent welfare effects.

The FOC with respect to frequency captures the optimization decision surrounding an additional flight; the marginal revenue recovered here, and the marginal cost, is of an additional flight. This second system of FOCs will be solved simultaneously with the slot constraint:

$$\frac{\partial \mathcal{L}_c}{\partial \lambda_c} = S_c - \sum_k f_k \tag{7}$$

I will separately estimate the MC_k of a flight (discussed below), and use it together with the demand-side parameters and data, to recover the shadow cost of an additional slot to a carrier, λ_c.

IV.2.2 Costs

In line with the transportation literature (McCarthy, 2001), transportation cost is increasing in distance flown, at a decreasing rate. A quadratic function, therefore, sufficiently captures the cost of fuel and flight crew that increases with distance, as well as the distance-invariant expenses of a flight. I use a simple function:

$$MC_k \equiv \kappa_0 + \kappa_1 dist_k + \kappa_2 dist_k^2 + \omega_k \tag{8}$$

where $dist_k$ is the actual distance flown by a flight and ω_k is a mean-zero error term. From the firm's FOC with respect to frequency in Equation 6, we have:

$$\widehat{MR_k} - MC_k - \lambda_c = 0$$

I substitute the functional form of the marginal cost to estimate the cost parameters:

$$\widehat{MR_k} = \kappa_0 + \kappa_1 dist_k + \kappa_2 dist_k^2 + \lambda_c + \omega_k \tag{9}$$

λ_c is identified using fixed effects for the carrier interacted with year.

IV.2.3 FOC with respect to price

I estimate the counterfactual vector of prices using the following FOC for pricing product i with respect to price:

$$\begin{aligned}
\frac{\partial \mathcal{L}_c}{\partial p_i} &= M\left(\frac{\partial s_i}{\partial p_i} p_i + s_i + \sum_{\forall j \neq i, j \in J_c} \frac{\partial s_j}{\partial p_i} p_j\right) = 0 \\
&\Rightarrow s_i + \sum_{\forall j \in J_c} \frac{\partial s_j}{\partial p_i} p_j = 0 \\
&\Rightarrow \widehat{s_i}(x_i, \widehat{p_i}, f_i) + \sum_{\forall j \in J_c} \frac{\partial \widehat{s_j}(x_j, \widehat{p_j}, f_j)}{\partial \widehat{p_i}} \widehat{p_j} = 0
\end{aligned} \tag{10}$$

where $\widehat{s_i}$ is the share of the i-th product (which is a function of the estimated price) calculated using Equation 2. This system of FOCs are solved for the price, p_i, under various counterfactual scenarios represented by different product characteristics, x_i's, and frequencies, f_i's. The system can be separated by market, since the optimal price only depends on within-market observables.

IV.3 Instruments

Table 6: First stage regression for instruments

	Price	Frequency
Number of rival passengers in direct legacy flights	0.0325***	
	(11.96)	
Number of rival passengers in direct low-cost flights	0.0214***	
	(3.97)	
Destination is hub	-10.72***	
	(-12.75)	
Total flights by the carrier from origin		169.1***
		(14.02)
Total nonstop flights by the carrier from origin		-0.831
		(-0.50)
Predicted total flights		-1.302*
		(-2.10)
Predicted total direct flights		7.307***
		(10.83)
Constant	281.8***	-6.964***
	(130.17)	(-13.44)
Observations	10,456	10,456
R-squared	0.0188	0.3916
F-value	66.91***	1682.05***

t statistics in parentheses

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

I assume that, in the short run, the existence of airlines, their hubs, networks, and routes are exogenous, but airlines can control the intensive margin – they choose how many flights to fly along a route, and the price of each seat. To that end, both market fare and flight frequency are endogenous, and need to be instrumented for.

Since I assume networks are exogenous, network features can instrument for flight frequency. Specifically, this includes instruments such as the total number of flights operated by the carrier out of the origin airport and the total number of nonstop flights from the origin airport operated by the carrier. Following Berry and Jia (2010), I also construct instruments for flight frequencies by first regressing segment departures on characteristics of the end cities, and then including the predicted segment departures as instruments.

To instrument for price using cost shifters that do not directly affect demand, I again exploit the assumption that network features (for example, other airport hubs) are exogenous to the Reagan market (Borenstein, 1989). For example, airlines have a larger pool of airplanes and flight crew available at hubs, and can thus better minimize non-weather delays for flights that use a hub as an endpoint.

I use as instruments for price: an indicator for whether the destination airport is a designated hub for the carrier, and the number of rival passengers traveling on direct flights to their destination (split by whether the carrier is legacy or low-cost).

IV.4 Identification

The type-specific logit model assumes that consumer tastes are drawn from a discrete distribution.[17] As such, I estimate $2k + 1$ parameters: a coefficient for each of the k product characteristics for leisure and for business travelers, and a parameter for the

[17] I assume two consumer types, so I have a bimodal discrete distribution to draw from. Assuming an underlying discrete distribution for consumer tastes also allow for analytic solutions for the inversion of share equations.

proportion of business travelers.

I begin with a stylized example, assuming that consumer tastes follow a simple logit model. In a market, there are two products identical in every aspect, except that one product is priced \$100 and the other \$500. If the market share is one-half for both the products, the introduction of a third product (identical to the first two) priced at \$100 will change the market share to one-third for each of the three products (readers might know this as the "Red-Bus Blue-Bus" problem (McFadden, 1974)). Our intuition, however, is that substitution is more likely to occur between the two \$100-products than between the products priced differently. If the observed market shares deviate from the results implied by the simple logit model, I can rationalize the observed market shares by assuming two coefficients on price. When comparing identical markets with different product assortments, the implied substitution patterns from observed market shares allow me to identify two distinct coefficients for each of the product characteristics.

The assumption that each of the coefficients and the proportion of business-travelers are the same across all markets allows me to uniquely identify each of the $2k + 1$ parameters. Allowing any of the parameters to vary across markets would deprive me of the across-market variation required to identify any of the parameters.

IV.5 Estimation

I recover the demand parameters $\theta_d = \{\alpha_r, \beta_r, \xi_{jt}, \gamma\}$ and supply parameters $\theta_s = \{\kappa, \lambda_c\}$ separately.

I begin with the demand estimation. In order to form the demand moment conditions, I first invert the share equation to solve for the unobservable characteristics, ξ_{jt}:

$$\xi_{jt} = s^{-1}(x_{jt}, p_{jt}, s_{jt}, \theta_d) \tag{11}$$

I find the vector of unobservables, ξ_{jt}, using the following contraction mapping:

$$\xi_{jt}^{H+1} = \xi_{jt}^{H} + \Big(\ln s_{jt}^{obs} - \ln s_{jt}(x_{jt}, p_{jt}, \xi_{jt}, \theta_d)\Big) \tag{12}$$

where, s_{jt}^{obs} is the observed share, s_{jt} is the market share estimated using Equation 2, and H is the H-th iteration.

My choice of using contraction mapping on the product unobservables, ξ_{jt}, follows Kalouptsidi (2012), and deviates from literature (Berry et al., 1995; Nevo, 2001) that solves for the mean utility, δ_{jt}, instead. Since a type-specific-coefficient logit model has multiple mean utilities, each corresponding to a consumer-type, I cannot invert the share equation to solve for a unique mean utility for a product. Kalouptsidi (2012) shows that theoretically, the two methods are equivalent.

The moment condition used in the estimation:

$$E\Big(h(z_t) \cdot \xi(x_t, p_t, s_t, \phi, \theta_d)\Big) = 0 \tag{13}$$

is formed by interacting my demand-side unobservables with a vector of demand-side instruments, z_t, the details of which are discussed in subsection IV.3.

I use the θ_d recovered here as inputs to obtain a better estimate of share, and iterate between the contraction mapping and moment condition to obtain a better estimation for θ_d. This process continues until the moment condition is minimized by a θ_d that returns a stable ξ_{jt}.

To estimate supply-side parameters, I use θ_d to compute the marginal revenue of a flight, $\widehat{MR_k}$, and back out the $\widehat{\kappa}$'s and the average shadow cost of a slot at Reagan Airport to an airline, $\widehat{\lambda_c}$.

To simulate a counterfactual world in which the merger is consummated without the divestment, I change the ownership of all slots belonging to American Airlines and US Airways to a new entity. In markets where both US Airways and American Airlines exist, I simulate a best-case scenario for the merged entity: where the best product unobservables between the two airlines (the highest ξ_{jt}) are assigned as the product unobservable of the new product, and a worst-case scenario: where the worst product unobservables between the two airlines (the lowest ξ_{jt}) are assigned as the product unobservable of the new product.

In order to compare the realized scenario with the best-case and worst-case scenario, I compute the expected consumer welfare (de Jong et al., 2005) in each scenario using the following equation:

$$\text{Expected Consumer Surplus} = \sum_r \gamma_r \frac{1}{\alpha_r} \ln \sum_{j=1}^{J} \exp \delta_{rjt} \tag{14}$$

Since I normalize the mean utility of the outside good to zero for all consumers, this expected consumer surplus is the average consumer surplus earned by a representative consumer in excess of the consumer surplus from the outside good.

V Results

V.1 Demand and Cost Coefficients

The demand coefficients in Table 7 are consistent with the two types in the model capturing leisure-traveler and business-traveler preferences. With this labeling, business-type travelers (price-elasticity of demand, $\varepsilon_D = -2.42$) are less price sensitive than leisure-type consumers ($\varepsilon_D = -4.56$), and dislike layovers more than leisure-type consumers do. Leisure travelers prefer going to tourist destinations more than business travelers, and prefer routes with higher price deviation (giving them the opportunity to purchase tickets from the cheaper end of the spectrum). I find the proportion of leisure travelers for routes involving Reagan Airport to be 53.42%, lower than the national average of 63% in 2006 estimated by Berry and Jia (2010).

Cost coefficients are estimated using Equation 9. The coefficients indicate that marginal cost is increasing at a decreasing rate for all relevant values of distance.[18]

V.2 Shadow value of slots

The estimated fixed effects for shadow values listed in Table 9 are from the third quarter of each year. The merger takes place between February 2013 – when the merger was announced, and April 2015 – when a single operating license was granted to the newly merged carrier. As such, we can think of 2012 prices to be pre-merger, and the 2015 prices to be post-merger-and-divestment.

Since I calculate the shadow value of the marginal slot using fixed effects, I capture the sum of the true shadow cost and any fixed costs not associated with distance (see

[18]Because Equation 9 is quadratic, the marginal cost is increasing in distance until distance equals 2,912 miles. However, this is already greater than the distance between Reagan Airport and Los Angeles, the airport farthest away from Reagan Airport in the contiguous United States.

Table 7: Demand Coefficients

Covariate	Type-Business	Type-Leisure
Constant	-0.5772	4.1593
	(1.2694)	(0.1076)
Layover	-5.3488	-4.9447
	(0.0911)	(0.0391)
Tourist	-0.1612	0.2332
	(0.0335)	(0.6695)
Logdist	-0.8545	0.5996
	(0.1274)	(0.0202)
Closest	0.0005	0.0010
	(0.00001)	(0.0021)
Legacy	-1.9311	-2.8259
	(0.2132)	(0.0507)
SD price	-0.0093	0.0634
	(0.0098)	(0.0507)
Price	-0.0104	-0.0195
	(0.00032)	(0.00056)
Frequency	0.0179	0.0072
	(0.00093)	(0.00030)
Proportion*	46.58%	53.42%
	(0.00000)	

Parentheses contain standard errors.

* The proportion of leisure-type travelers is the complement of proportion of business-type travelers, and thus do not have standard errors of its own.

Equation 9). Mathematically,

$$\widehat{FE_c} = \lambda_c + \kappa_0 \tag{15}$$

where $\widehat{FE_c}$ is the estimate using fixed effects, λ_c is the true shadow value for airline c, and κ_0 is the fixed cost of operating a flight. We can estimate bounds on the true shadow value by making extreme assumptions about the fixed cost κ_0 and the true shadow value.

Assumption 1: The shadow value of the dominant incumbent (US Airways/American Airlines) is zero.

Table 8: Cost Coefficients

Covariate	Coefficient
Miles flown (squared)	-0.0013
	(0.0001)
Miles flown	7.6165
	(0.4205)

Parentheses contain standard errors.

This assumption can be justified by noting that US Airways' slot utilization (calculated by the number of flights scheduled as a percentage of total slots) is the lowest among all carriers, hovering around the minimum required 80% mark.

A corollary of this assumption is that the estimated fixed effect for US Airways, $\widehat{FE_{cUS}}$ from Equation 15 is equal to κ_0. Deducting this estimated κ_0 from all estimated fixed effects gives us the lower bound on true marginal shadow value of slots for all airlines.

Assumption 2: The fixed cost of operating the marginal flight is zero.

This assumption can be justified by arguing that all fixed costs associated with operating the marginal flight are captured by a quadratic function of distance. Since doubling the distance will cause the cost of operating the marginal flight to less than double, this can capture the distance-invariant costs like gate agents' wages and jet bridge rentals when distance is zero.

We can explain the low slot utilization with the inefficiencies of a hub-and-spoke network when capacity is constrained; slot utilization is lower for all legacy carriers, lending credence to this explanation.

A corollary of this assumption is that $\widehat{FE}_c = \lambda_c$ for all airlines, since $\kappa_0 = 0$, which establishes an upper bound on the true marginal shadow value of slots for all airlines.

The two extreme assumptions, combined, give us bounds on the true shadow value of the marginal slot to an airline. The true shadow values implied by Assumptions 1 and 2 are outlined in Table 9:

Table 9: (Bounds on) shadow values implied by estimated fixed effects, in dollars

	2012 (Before merger)			2015 (After merger)		
	Slots	Lower bound	Upper bound	Slots	Lower bound	Upper bound
US Airways	473	0.00	152.62	491	0.00	233.79
American	118	272.46	425.08			
Delta	104	353.03	505.65	104	311.76	545.55
United	82	207.81	360.43	82	240.87	474.65
Southwest	—	—	—	82	617.93	851.72
JetBlue	20	862.92	1015.54	60	685.51	919.30

The shadow value of slots reflects the worth, in dollars, of the last slot controlled by an airline. We observe two patterns. Firstly, the value of slots increases for most airlines (except the merged entity), even those not directly involved with the divestment (United and Delta, although the lower bound estimate for Delta falls post-merger). This perhaps reflects a general upward trend in the economy, making each flight more valuable than in the past, or a collective increase in the shadow value of each slot due to a possible reduction in competition on legacy routes.

Second, the shadow value of a slot is decreasing in total slot allocation. This is best observed in Figure 1, with the slope of a fitted line through the points being -1.24 (p=*0.017*). A corollary of this is that the divested slots were allocated to firms who valued the slots the most (JetBlue and Southwest). If a regulator offered an

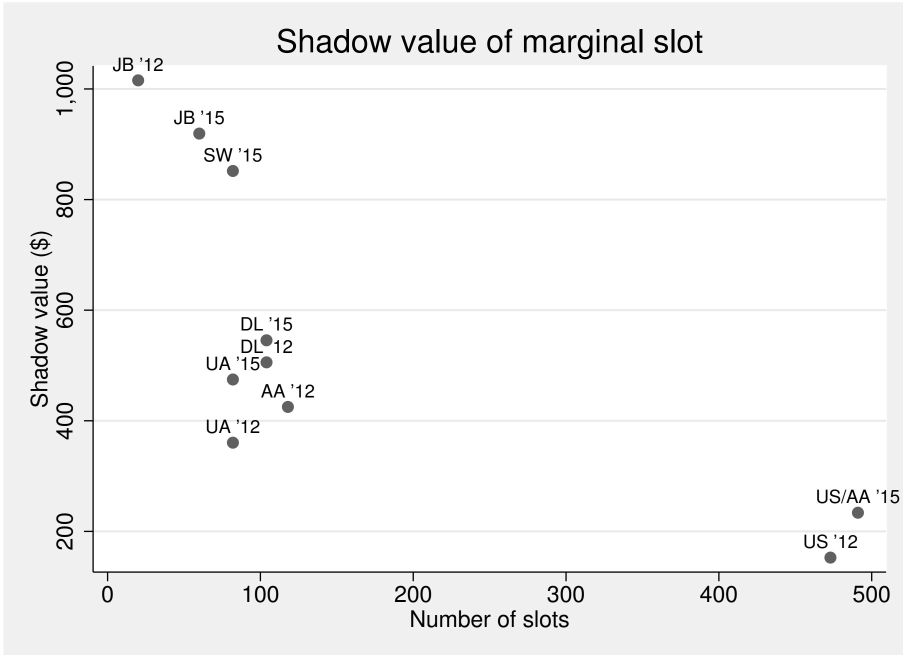

Figure 1: Shadow value of the marginal slot against total slot allocation (upper bound)

additional slot to a carrier, with the alternate being that the slot is never created, the shadow value represents the amount of money an airline is willing to spend to buy this additional slot from the regulator. This is not the willingness-to-pay for an additional slot if the alternate is another airline buying the slot. Airlines might be incentivized to foreclose their competitor from obtaining the slot,[19] which is not reflected in the values in Table 9. Since airline products are substitutes, the shadow values in Table 9 reflect the minimum price an airline is willing to pay for an additional slot at a competitive auction (where the alternate is a rival carrier buying the slot).

Lastly, we can engage in some back-of-the-envelope calculation to externally verify the shadow values of the slots. We know that American Airlines received $425

[19]The foreclosing incentive is as follows: Delta Airlines values an additional slot at $546, but if Delta doesn't buy a slot for sale and it is bought by United Airlines, this additional slot could be used by United to directly compete with Delta on a specific route, reducing Delta's profits by, say, $200. This means that when the alternate to buying a slot is a rival carrier obtaining it, Delta is willing to pay up to $746, more than the shadow values in Table 9, in order to *foreclose* their rival.

million for the mandatory sale of 138 slots (Maxon, 2014), or \$3.08 million per slot. Using US/American's post-merger upper bound, the transaction price implies that the merged entity expected to recoup the value of the slot in 36 years (\$3.08 million ÷ 365 days per year ÷ \$234 per day). Using JetBlue's post-merger lower bound, the transaction price implies that JetBlue expects to recoup the value of the slot in 12 years (\$3.08 million ÷ 365 days per year ÷ \$686 per day). Put differently, the seller's upper bound is less than the buyer's lower bound on the shadow value, which is reasonable.

V.3 Counterfactual Simulations

I consider two counterfactual studies. In both studies, I investigate the effects of the merger without any accompanying divestment by assigning all slots owned by US Airways and American Airlines to a newly merged entity, and simulating prices for all the products at the pre-merger level of frequency using Equation 10. Since the suite of products available following this counterfactual merger does not exist, I need to make assumptions about the unobservable product characteristics of the merged entity, ξ_{jt}. The choice of these assumptions distinguishes the two counterfactuals.

For the first counterfactual, I consider the best-case scenario for the merged entity. For markets involving a US Airways or American Airlines product, I take the best product unobservables (the highest ξ_{jt}) among all products offered by this newly merged entity, and ascribe that to all products being offered by this merged firm in this market. This simulates positive synergy following the merger. A second counterfactual ascribes the worst product unobservables (the lowest ξ_{jt}) to all the products offered by the merged firm in the market, simulating a worst-case scenario where the merger causes a fall in unobserved product quality.

Table 10: Consumer Surplus under Counterfactual Scenarios

	(1) Before merger	(2) Observed Merger w/ Divestment	(3) Best-case for US/AA	(4) Worst-case for US/AA
Median price ($)	290.95	267.65 (-8.01%)	280.70 (-3.52%)	312.90 (7.54%)
Total passengers (millions)	122.168	132.474 (8.44%)	124.014 (1.51%)	125.526 (1.93%)
Average Expected Consumer Surplus per Passenger ($)	27.93	35.05 (25.49%)	28.44 (1.81%)	27.69 (-0.86%)
Total Consumer Surplus ($ millions)	3.574	4.485 (25.49%)	3.630 (1.56%)	3.562 (-0.32%)

Parentheses contain percentage change from pre-merger levels

For each of the counterfactual studies, I compute the expected consumer surplus of a representative consumer, as well as the total consumer surplus. Table 10 shows that the divestment brought gains in both the average consumer surplus per passenger as well as the overall consumer surplus. A merger without divestment engenders negligible change in consumer surplus; since US Airways and American Airlines had no duopoly routes before the merger, the upward pricing pressure following the merger is mitigated by competition from other firms.

To understand the increase in average consumer surplus, I look at the distribution of gain in consumer surplus (difference between Table 10, Columns 2 and 1) across each market, weighted by the number of passengers in the market. Figure 2 shows that most consumers make modest gains in consumer surplus; the right tail of the distribution is dominated by routes that experience entry by low-cost carriers, which depresses price and increases quantity of travelers. The distribution of consumer surplus gains for business travelers is visually similar.

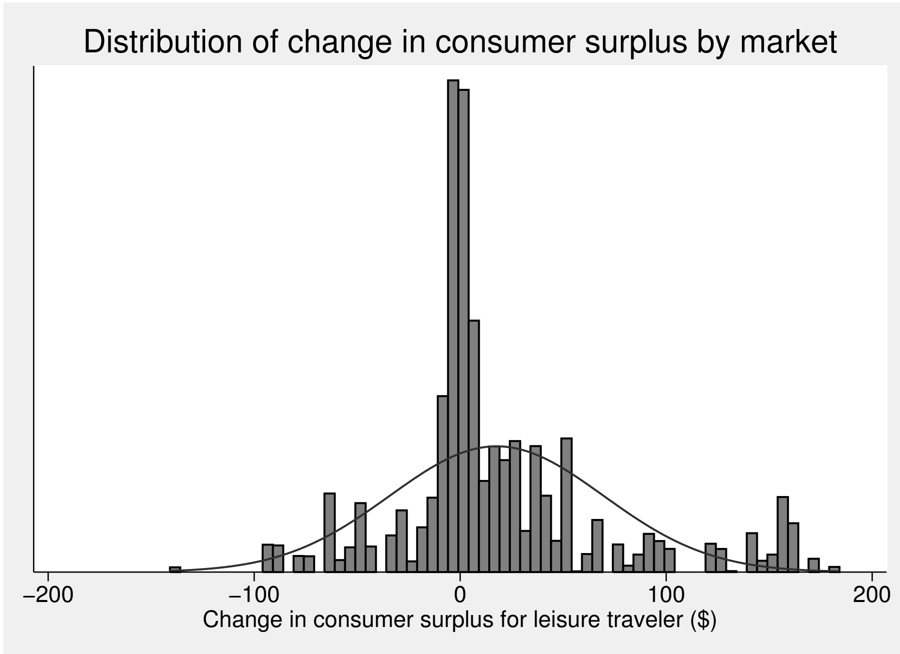

Figure 2: Distribution of change in consumer surplus by market

I decompose this increase in consumer surplus by unpacking the components of the utility function. The divestment mechanically increases the proportion of flights being operated by low-cost carriers; there is a 7.6% drop in the amount of legacy carrier products. Low-cost carriers like Southwest and JetBlue mostly offer point-to-point service, which increases the number of products that are direct flights; the proportion of flights with connections dropped by 17% following the divestment. Lastly, legacy carriers have a lower slot utilization than low-cost carriers – this may be due to scheduling considerations due to their hub-and-spoke network. As a result, divesting the slots to low-cost carriers increases slot utilization, and as a result, number of flights on a route. Passengers dislike layovers, dislike legacy carriers, and prefer higher frequency (Table 7). This results in a higher mean utility of a product, and therefore, higher expected consumer surplus for an individual consumer. This effect can be offset by the increase in price following the merger; in the scenario involving divestment (Table 10, Column 2), the price in fact drops due to increased competition by low-cost carriers, causing an increase in expected consumer surplus per person. Since the channel for increase in consumer welfare stems from the divestment, the counterfactual simulations (Table 10, Columns 3 and 4) show much smaller change for both the expected consumer surplus per person and the total consumer surplus.

The increase in mean utility due to the divestment causes more passengers to fly than before (Table 10, Column 2). The gain in number of passengers, coupled with an increase in the expected consumer surplus for the representative consumer, causes an increase in total consumer surplus of about $911,000 (25.49%) as a result of the divestment.

Lastly, I find a weak but positive correlation between market size and marginal revenue (a 10,000-person increase in market size increases marginal revenue by $7, significant at a 99% level). This indicates that smaller markets are less profitable, at least at the margin. If the newly merged entity were to drop flights without any restrictions, this implies that they are likelier to drop flights to a smaller market than to a larger one. Assuming that service to small markets have positive externalities, I find that the separate settlement (United States v. US Airways Group, 2014) between the DOJ and the newly merged entity (where the latter promised not to terminate any routes serving small communities) mitigated this effect.

VI Conclusion

In this paper, I suggest a method to incorporate structural remedies into the standard demand and supply model. While the shadow value of a slot is implicit in all airline literature, I show how to separately identify and put bounds on the marginal value of a slot for an airline. I find that the marginal value is decreasing in slot endowment, and is higher for low-cost carriers than for legacy carriers. Airlines with high slot endowment must balance gains from higher frequency with loss of inframarginal revenue from increase in quantity, resulting in high-endowment airlines putting a lower valuation on their marginal slot compared to a low-endowment airline.

This validates the DOJ decision to redistribute the divested slots to new low-cost entrants (like Southwest and JetBlue) as opposed to rival legacy incumbents. My retroactive analysis of the US Airways-American Airlines merger and the structural remedies imposed by the DOJ confirms that the divestment increased the average consumer surplus of a passenger traveling in the Reagan market by 25.49%, or an increase of $7.12 per passenger per one-way flight. There are three main reasons for this gain in consumer surplus: (a) the divestment allocated slots to low-cost carriers, (b) LCCs have more point-to-point direct service, and (c) LCCs have higher slot utilization. A counterfactual simulation (holding frequency constant) where the merger is consummated without divestment shows that gains in consumer surplus are negligible, and depend on assumptions about product unobservables; this shows that the divestment was the main driver for gains in consumer welfare. However, Figure 2 shows that while the average consumer surplus did increase, this gain is unevenly distributed across markets. In future work, I intend to investigate the correlation between market characteristics and gain (or loss) in consumer surplus.

My approach can be generalized to predict the impact of any future proposal that involves a reallocation of slots, whether voluntary or mandated by a regulator. My empirical approach can be extended to estimate a willingness-to-pay for a slot, which will include the shadow value to the carrier plus the gains from foreclosing a rival carrier. In this paper, I show the importance and efficacy of using airport slots (and gates in general) in crafting antitrust policy for airlines, and suggest an empirical approach to evaluate the impact of such antitrust policies.

References

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

Berry, S., Carnall, M., and Spiller, P. (2006). Airline hubbing: Costs and demand. In Lee, D., editor, *Advances in Airline Economics: Competition Policy and Antitrust*, volume 1. Elsevier.

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

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, pages 841–890.

Berry, S. and Waldfogel, J. (2010). Product quality and market size. *The Journal of Industrial Economics*, 58(1):1–31.

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

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

Cabral, L. M. (2003). Horizontal mergers with free-entry: why cost efficiencies may be a weak defense and asset sales a poor remedy. *International Journal of Industrial Organization*, 21(5):607–623.

Ciliberto, F., Murry, C., and Tamer, E. T. (2016). Market structure and competition in airline markets. *Available at SSRN 2777820*.

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

de Jong, G., Pieters, M., Daly, A., Graafland-Essers, I., Kroes, E., and Koopmans, C. (2005). Using the logsum as an evaluation measure: Literature and case study. *RAND Corporation*.

Dobson, P. W. and Piga, C. A. (2013). The impact of mergers on fares structure: Evidence from european low-cost airlines. *Economic Inquiry*, 51(2):1196–1217.

Hotelling, H. (1929). Stability in competition. *The Economic Journal*, 39(153):41–57.

Kalouptsidi, M. (2012). From market shares to consumer types: Duality in differentiated product demand estimation. *Journal of Applied Econometrics*, 27(3):333–342.

Kim, E. H. and Singal, V. (1993). Mergers and market power: Evidence from the airline industry. *The American Economic Review*, pages 549–569.

Kwoka, Jr., J. E. (2013). Does merger control work? A retrospective on US enforcement actions and merger outcomes. *Antitrust Law Journal*, 78.

Maxon, T. (2014). American Airlines sells New York, Washington slots for more than $425 million. *Dallas News*.

McCarthy, P. (2001). *Transportation Economics: Theory and Practice: A Case Study Approach*. Wiley.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In Zarembka, P., editor, *Frontiers in Econometrics*. New York: Academic Press.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2):307–342.

Peters, C. (2006). Evaluating the performance of merger simulation: Evidence from the US airline industry. *The Journal of Law & Economics*, 49(2):627–649.

Polsby, D. R. (2001). Airport pricing of aircraft takeoff and landing slots: An economic critique of federal regulatory policy. *California Law Review*, pages 779–816.

Reitzes, J. D., McVeigh, B., Powers, N., and Moy, S. (2015). Competitive effects of exchanges or sales of airport landing slots. *Review of Industrial Organization*, 46(2):95–125.

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

Tenn, S. and Yun, J. M. (2011). The success of divestitures in merger enforcement: Evidence from the j&j–pfizer transaction. *International journal of industrial organization*, 29(2):273–282.

United States v. US Airways Group (2014). *United States District Court for the District of Columbia*, 13-cv-1236.

United States v. US Airways Group & AMR Corporation (2013). *United States District Court for the District of Columbia*, 13-cv-1236.

Vasconcelos, H. (2010). Efficiency gains and structural remedies in merger control. *The Journal of Industrial Economics*, 58(4):742–766.

“Use It or Lose It”, or “Cheat and Keep”? Effects of Slot Restrictions on Airline Incentives

Ratib M Ali and Vera Sharunova

April 12, 2022

Abstract

The Federal Aviation Administration manages congestion in high-density airports by capping the number of flights permitted in any given hour and allocating the rights (or slots) to a takeoff or landing among airlines. Airlines must use their slots at least 80% of the time to keep them for the next season. This rule creates a perverse incentive for airlines to hold on to underutilized slots by operating unprofitable flights instead of forfeiting these slots to a rival. Using exogenous removal of slot control at the Newark Airport in 2016, we investigate the lengths at which airlines go to meet the minimum requirements that let them keep the slots while violating what a neutral observer might call the “spirit” of the regulation.

I Introduction

Air traffic congestion has been responsible for a significant welfare loss in the US economy. Peterson et al. (2013) estimate that a 10% reduction in flight delays[1] would increase net total welfare by $17.6 billion. The congestion problem has been a major concern for all stakeholders since the 1960s and is exacerbated by limited airport capacities.

Regulators all over the world, including the Federal Aviation Administration (FAA) in the US, manage congestion by capping the number of hourly flights and rationing the rights to perform a take-off or a landing within a given timeframe (also known as "slots") among airlines. Slot controls are currently in place in most major airports worldwide and in three US airports – JFK and LaGuardia Airports in New York City, and Reagan National Airport near Washington, DC.

Following the International Air Transport Association (IATA) procedures, the FAA revisits slot allocations each year at the start of the winter and summer seasons. Airlines are allowed to keep their slot holdings for the next season provided that they comply with the use-it-or-lose-it rule by using their slots at least 80% of the time during the current season. This approach biases the allocation process in favor of legacy carriers, who were incumbents when slot control was first introduced in 1969 and still hold the majority of slots in slot-controlled airports.[2]

While this regulation has been successful in managing congestion and delays, it may also create an incentive for airlines to hold onto unprofitable slots in order to

[1]Flight delays are measured as a fraction flights that are delayed by 15 minutes or more.
[2]Calculated from the slot holder reports for the winter season of 2018 published by the FAA.

keep their competitors out of highly-demanded airports (GAO, 2012).[3] Specifically, when slots at an airport are limited, airlines not only want a slot to operate a flight on a given route, but also to prevent a rival from controlling said slot and competing directly against the incumbent, possibly on other, more profitable, routes. The incentive to foreclose manifests in slot burning – using slots at a loss instead of forfeiting them to a rival and incurring a greater loss in profits.

Given the prevalence of slot control, it is important to study its effects on consumers and investigate the alleged anticompetitive incentives it creates for airlines. This paper uses reduced-form analyses to assess evidence of slot burning, relying on a natural experiment created by the removal of slot control at Newark Airport in November of 2016. Removal of slot control eliminates historical precedence created by the use-it-or-lose-it rule and allows entry of new airlines who did not previously hold slots at the airport, thereby eliminating slot-burning incentives.

We measure the extent of slot burning by using the empirical probability of observing a small flight on a given route as a proxy variable. Anecdotally, we have heard of airlines operating frequent flights using smaller aircraft to use up multiple slots as opposed to carrying the same number of passengers in fewer flights using larger aircraft (GAO, 2012). Slot incumbents, owing to their large slot endowments, may be more prone to slot burning. In line with the anecdotes, we find circumstantial evidence for airlines burning slots: frequency of small flights between Newark and Philadelphia Airports, only 80 miles apart, decreases by nearly 77% following removal of Newark's slot control in 2016, with slot incumbents accounting for 35% of the drop.

[3] Due to the disruptive effects of the COVID-19 pandemic, the FAA suspended the use-it-or-lose-it rule on March 11, 2020 in order to relieve airlines from the need of flying 'ghost planes', in other words, slot burning (Pallini, 2020).

Indeed, we find that slot incumbents are twice as likely to operate smaller aircraft if slot restrictions are in place. We note that slot burning occurs most often during relatively offpeak hours of 10am to 1pm – a timeframe where there is enough demand to support some flights, but not enough demand to warrant usage of larger aircraft.

Airlines, however, contend that using multiple smaller flights is needed to meet demand for schedule-sensitive passengers. We address this argument by looking at the probability of flying small aircraft to the same destination within a 30-minute timeframe. We find that slot incumbents are 75% more likely to fly consecutive small flights when Newark Airport is slot-controlled.

We also find that slot restrictions are associated with increased airfares and decreased delays, suggesting that, while slot controls are effective in decreasing congestion, they do so at the expense of higher airfares due to restricted entry. In addition, we analyze changes in several quantity metrics, namely, the number of seats(-miles) offered on a route (as a proxy for quantity supplied) and the number of passengers(-miles) transported (as a proxy for quantity demanded). We find that slot incumbents offer more seats and transport more passengers under slot restrictions. Together with the increase in airfares, this observation suggests that both demand and supply shift under slot control. We discuss a potential rationale for these shifts and their effects on consumer surplus in Section V.

Lastly, we investigate the patterns of entry to and exit from markets, finding that slot liberalization resulted in entry by low-cost carriers to certain routes. We highlight the competing interests that must be met when aviation authorities consider mitigating congestion through slot restrictions; namely, effects on price, delays, entry, and the incentives of slot incumbents to burn slots. An evaluation of consumer welfare, as a function of frequency, prices, and delays, is left for future work.

The rest of the paper is organized as follows: Section II surveys the relevant literature. Section III reviews the industry background and introduces the data. Section IV develops a theoretical model that informs our reduced-form analyses. Section V presents our hypotheses and discusses the results. Section VI provides an overview of the entry and exit decisions following Newark's reclassification. Section VII concludes.

II Related Literature

Our research contributes to the literature on airlines' access to airport facilities and its effect on downstream market outcomes, specifically airfares and congestion. Ciliberto and Williams (2010) investigate whether access to airport gates, which is usually determined by long-term exclusive contracts between airlines and airports, allows airlines to charge higher prices on flights in and out of their hubs. In the first stage, Ciliberto and Williams (2010) recover carrier-route fixed effects from a reduced-form pricing equation as a measure of the hub premium, which they later regress on measures of access to airport gates. They find that an increase in the percentage of gates controlled by an airline is associated with an increase in its airfares, especially in more congested airports as defined by the number of departures per gate.

Snider and Williams (2015) study the changes in airfares following the change in access to airport facilities mandated by the AIR-21 Act. They use a regression discontinuity approach to identify changes in airfares at airports that were required to improve access to their facilities, relative to airports that were exempt from this requirement based on two threshold rules. They find that airfares decrease by 13.4% on routes where one airport is covered by the legislation and by 20.2% on routes where both end points are covered.

In a recent paper, Fukui (2019), similarly to us, uses the exogenous change in slot restrictions at Newark Airport in 2016 to study the impact of slot control on average airfares. They use a difference-in-differences approach, where the treatment group consists of routes to or from Newark and the control group consists of routes to or from the two other NYC airports – the JFK and LaGuardia Airports. They find that the average fare on Newark routes decreases by about 2.5% relative to the JFK and LGA routes, with the majority of the effect coming from non-dominant Newark airlines. Our study documents the effect on airfares by using all major airports as the control group, as opposed to just JFK and LGA, although our primary focus is on slot burning.

In another study from Newark Airport, Luttmann (2019) exploits reinstitution of slot control at the JFK and Newark Airports in 2008 to evaluate the effectiveness of slot control in managing delays. Using the 2007-2008 data, Luttman (2019) finds no evidence of a reduction in delays at both airports. They suggest that these findings are consistent with the internalization hypothesis claiming that delays at an airport decrease with an emergence of a dominant airline. In contrast to Luttman's results, we find that delays at Newark Airport increase following the removal of slot restrictions in 2016. Figure 1 shows that the average length of delays at Newark Airport diverges away from other NYC Airports following the abolition of slot restrictions. Additionally, we find a decrease in delays following the 2008 classification as well (unlike Luttman, 2019), but we focus on the 2016 (as opposed to the 2008) reclassification to abstract away from the confounding effects of the Great Recession.[4]

A 2012 report on slot-control rules published by the Government Accountability Office is the closest to our paper and has inspired our proxy variable for slot burning. In particular, GAO proposed three indicators of airline slot burning: (i) using smaller aircraft; (ii) flying to the same destination at higher daily frequencies; and/or (iii) operating flights with lower average load factors (passenger-to-capacity ratio). In particular, GAO compares the number and proportion of small aircraft flights (under 100 passengers) to and from slot-controlled airports to those from other large domestic hubs, controlling for flight distance and other relevant characteristics. They find that the odds that a flight to and from a slot-controlled airport uses a small aircraft are 75% higher than the odds for a flight to and from other large hub airports that are not slot-controlled (GAO, 2012). The evidence that GAO finds is only suggestive since the estimates are not causal and rely on how well other large domestic airports act as a control group for the slot-controlled airports. We improve upon the GAO's methodology by considering a natural experiment created by exogenous removal of slot restrictions at Newark Airport in 2016.

A related paper looks at the divestment of slots as a structural remedy to the 2013 US Airways-American Airlines merger to back out airline response to the reallocation of slots designed to promote competition and the associated effects on consumers (Ali, 2020). This paper, instead, summarizes airline behavior and its effects on consumers following a wholesale abolition of slot controls.

Lastly, Swaroop et al. (2011) investigate whether more US airports need slot

[4]We submitted FOIA requests to PANYNJ asking for Newark's time-stamped departure and arrival data from 2007-2009, but, after a diligent search, the agency could not locate any responsive records.

control and if the existing slot levels are optimal. In particular, they use an econometric model to quantify the costs of schedule delay, i.e. costs that a passenger suffers from having to choose the departure time from the set airline schedule as opposed to flying at her preferred time, and the costs of queuing delay, resulting from congestion. They later simulate optimal slot control policies at major US airports that minimize the sum of schedule and queuing delay costs and conclude that slot control should be implemented at 12 additional airports and slot caps decreased at the already slot-controlled airports. However, they do not take into account the impact of likely airfare increases as a result of constrained capacity and do not take into account anti-competitive effects generated by slot burning.

III Industry Background and Data

III.1 Slot Control at Newark Airport

The history of slot control in the United States goes back to the introduction of the High Density Rule (HDR) in 1969. The HDR capped the number of hourly arrivals and departures at five major airports – JFK, LaGuardia, Newark, O'Hare, and the Reagan Airport – and was seen at the time as a temporary measure to curb growing delays. The rule was suspended in Newark a year later since the number of flights was well under the cap, even at peak times. However, the HDR proved to be successful in managing congestion in the rest of the airports, so the FAA extended it indefinitely in 1973.

Initially, the slots were allocated by a group of airline representatives, the so-called scheduling committees, on a voluntary concession basis. However, after the deregulation of the airline industry in 1978, scheduling committees had been having difficulties agreeing on slot allocations, and the antitrust immunity which made their existence possible invited scrutiny from the FTC. By 1986, the FAA replaced scheduling committees with the slot-allocation procedures as they are currently known. In particular, the FAA instituted the use-it-ot-lose-it rule that stipulated withdrawal of slots that did not clear the 65% usage threshold. The minimum utilization threshold was increased to its current level of 80% in 1992.[5]

Since then, several government agencies[6] studied the effects of the High Density Rule on the quality of air service and concluded that the HDR was limiting competition and preventing improvements in the quality of service, partly because new entrants could not enter the slot-controlled airports. Eventually, the AIR-21 Act (2000)[7] called for a phase-out of the HDR at Chicago O'Hare Airport by July 2002 and at JFK and LaGuardia Airports by January 2007. After the expiration of slot control rules, air carriers have promptly increased their operations in JFK, making 2007 one of the worst years in terms of delays. All three of the New York City metropolitan area airports were affected by congestion at JFK (see Figure 1), so the FAA temporarily reinstated slot control soon after. Even though Newark remained slot-free from 1970 to 2008, the FAA decided to preemptively institute slot control rules at Newark as well, fearing that air carriers would shift their operations from JFK and LaGuardia and create additional congestion at Newark.[8]

Generally, existence of slot control at an airport entails: (i) caps on the number

[5] Amdt. 93-65, 57 FR 37315, Aug. 18, 1992.

[6] GAO, Airline Competition: Industry Operating and Marketing Practices Limit Market Entry, GAO/RCED-90-147 (Washington, D.C.: Aug. 29, 1990); National Research Council Transportation Research Board, Entry and Competition in the U.S. Airline Industry: Issues and Opportunities, Special Report 255 (Washington, D.C.: 1999); Department of Transportation, Study of the High Density Rule: Report to Congress (Washington, D.C.: May 1995).

[7] The Wendell H.Ford Aviation Investment and Reform Act for the 21st Century. Public Law 106-181.

[8] 73 FR 60543.

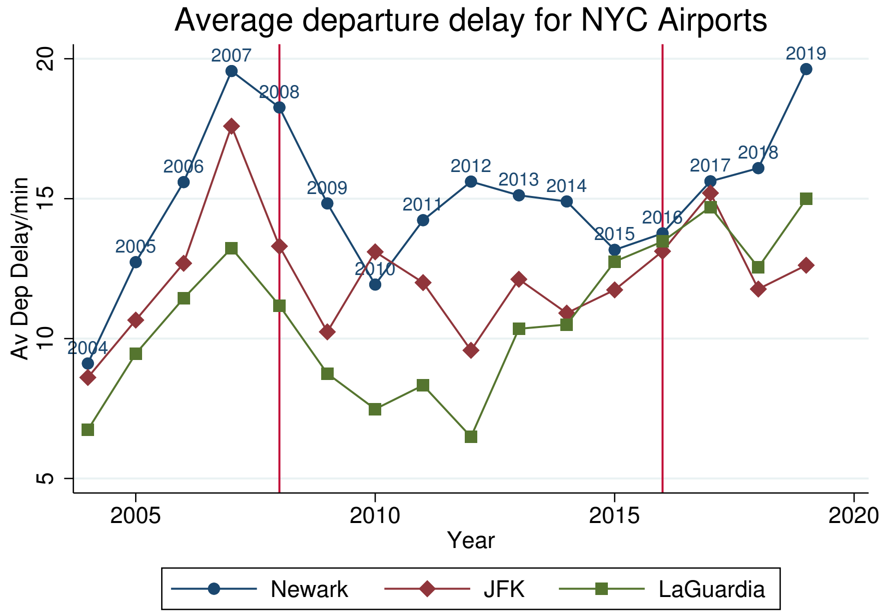

Figure 1: Average departure delays at NYC airports, in minutes.

of arrivals and departures performed within a 30-minute and one-hour timeframe; (ii) minimum usage requirement applied over a pool of slots within a given timeframe over a predefined period (henceforth, referred to as a slot period; differs across slot-controlled airports). For Newark Airport, the limit on the number of operations was set at 44 within each 30-minute window and 81 within each one-hour window from 6am to 10:59pm every day of the week.[9] The compliance with the minimum usage requirement was thus determined for each day of the week within a 30-minute and one-hour time periods, for example, Mondays from 6:00 to 6:30am during the winter season of 2015.

Newark's slot control rules were in place until the winter season of 2016. According to the FAA, the reasons for removal of slot restrictions in 2016 at Newark were three-fold: (i) improved capacity at JFK, following the runway reconstruction scheduled to begin in 2017, was expected to decrease the spillover effect that prompted slot control

[9]14 CFR 93.163(b) of January 1, 2009.

back in 2008; (ii) improvement in on-time performance and duration of delays from 2007 to 2015, and (iii) the FAA's prediction of future demand and capacity at Newark Airport suggested that the slot restrictions were no longer necessary.

As for the first reason, both JFK and LaGuardia underwent runway reconstructions that temporarily reduced the airports' capacities in 2017. The total number of scheduled flights at JFK and LGA decreased by about 3,000 and 1,200 respectively, relative to the 2016 levels, reversing the historical trends. The number of scheduled flights promptly returned to and exceeded the 2016 levels in 2018, after the construction projects were completed. In light of these events, we investigate if the FAA removed slot control at Newark to allow air carriers to reschedule their operations from JFK and LGA to Newark in anticipation of restricted capacity. We find no evidence of shifts in JFK's operations. However, we do find that the LGA routes that experienced a decrease in the number of scheduled flights in 2017 or 2018 (relative to 2016) tend to experience an increase in scheduled frequency at Newark. Routes that potentially shifted from LGA to Newark are not a part of the sample of airports we use to test for slot burning, therefore we believe that possible shifts in operations did not affect patterns of aircraft usage at Newark in any spurious manner.[10]

The second and third reasons cited by the FAA are both predicated on current improvements in on-time performance, as well as a belief held by the FAA (based off their future demand prediction at Newark Airport) that relaxing slot restrictions will not result in congestion due to entry. However, government and industry reports, as well as this study, show an increase in flight delays following the 2016 reclassification. This implies that any gains in on-time performance between 2009 and 2015 were likely due to effectiveness of slot-control restrictions. Figure 1 above shows that the average

[10]See Appendix B for more details.

delay for flights departing from Newark Airport fell in 2008 when the airport was escalated to a slot-restricted airport, but the average delays went back up to its pre-2008 levels following a de-escalation to a slot-facilitated airport. No such reversions occurred for other NYC airports in 2016 that kept slot controls and completed runway reconstructions by 2018. As a result, we believe that the reclassification was not endogenous to any sustained or systemic changes in schedule management at Newark Airport, and therefore, can exploit the reclassification as an exogenous change.

III.2 Suggestive Evidence of Slot Burning at Newark Airport

Interestingly, the FAA's review of Newark's operational performance concluded that scheduled demand was consistently below the 81 hourly flight cap, yet the FAA could not accommodate requests for new flights in summer of 2016 as the allocated slots reached the limit.[11] This conclusion suggests that the incumbent carriers might have relied on slot burning to meet the use-it-or-lose-it requirements and keep new entrants out of the airport.

As mentioned previously, our proxy variable for slot burning is usage of small aircraft. If an airline is burning a slot in order to prevent a competitor from acquiring it, the airline would minimize losses associated with operating an unprofitable flight by flying a smaller aircraft. To this end, Figure 2 highlights three stylized facts. First, slot-controlled airports use more small aircraft than non-slot-controlled airports. Second, in line with the FAA's evaluation of capacity usage at Newark prior to 2016, the share of small flights in Newark is around three times higher than in the other slot-controlled airports. Third, there is a meaningful change in the usage of small aircraft in Newark

[11] From 81 FR 19861: "For example, in the 3 p.m. through 8:59 p.m. local hours, weekday scheduled demand in the May-August period averaged 71 flights per hour in 2011, 74 flights per hour in 2013, and 72 flights per hour in 2015. [...] At the same time, the FAA denied requests for new flights as slots are allocated up to the scheduling limits."

Airport around 2016 that cannot be explained by a general time trend.

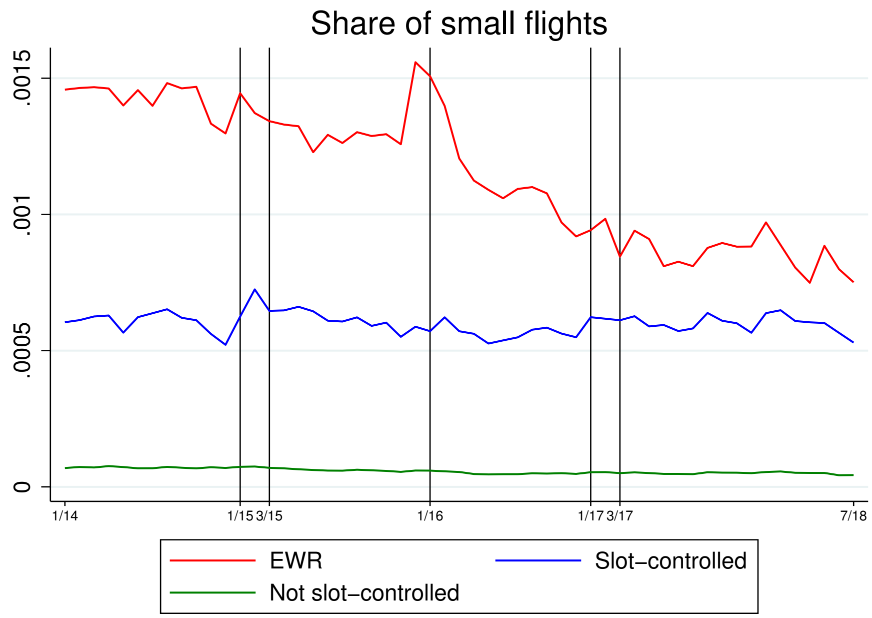

Figure 2: Percentage of small flights at Newark and other top-28 airports. We use data from the first quarter of 2015 (pre-reclassification) and 2017 (post-reclassification) in our analyses.

Moreover, airlines with relatively large pools of slots under their control may burn slots more often relative to airlines with smaller slot endowments. We refer to such airlines as slot incumbents, and we split our analysis by whether an airline is a slot incumbent or not. Table 1 below summarizes the total number of daily slots held by each airline at Newark in 2015. United, together with its regional partners, like Air Wisconsin and Republic Airways, held 869 daily slots, nearly 80% of all slots available at Newark.

In addition to varying daily slot endowments across airlines, the same airline generally holds a different number of slots in each slot period. This is due to the fact that slots granted to airlines are tied to a particular one-hour slot period.

Airline	Daily slots	Percentage
American	64	5.87
Alaska	4	0.37
Delta	65	5.96
JetBlue	40	3.67
Southwest	35	3.21
United	869	79.72
Virgin	13	1.19

Table 1: Daily slot holders, 2015.

Furthermore, compliance with the use-it-or-lose-it requirements is determined based on utilization of slots in the same 30-minute or one-hour slot period during a scheduling season. These two facts imply that some slot periods may experience more slot burning than others due to lack of demand or differing slot-burning incentives of airlines. Taking this observation into account, we revisit patterns of usage of small aircraft by time of arrival to/departure from Newark Airport. Figure 3 suggests that, under the slot control regime, the empirical probability of observing a small flight within the 10am-1pm timeframe is higher relative to the 4-8pm timeframe, while it is relatively uniform after slot control rules are lifted. We employ a Kolmogorov-Smirnov test for equality of the two distributions and reject the null hypothesis that the two distributions are the same with a D-value of 0.0468 (p=*0.000*).

We also investigate in which slot periods United, the slot incumbent, and the rest of Newark's airlines may be burning slots. As discussed more rigorously in section IV, we incorporate slot and minimum usage constraints into airlines' profit-maximization problems. Whenever an airline's minimum usage constraint is binding and the slot constraint is slack, we interpret such a phenomenon as slot burning. Figures 4 and 5 attempt to assess if United and Newark's low-cost carriers complied with said constraints. The top red line represents an airline's slot capacity – the average number of slots available to United or the group of low cost carriers in a

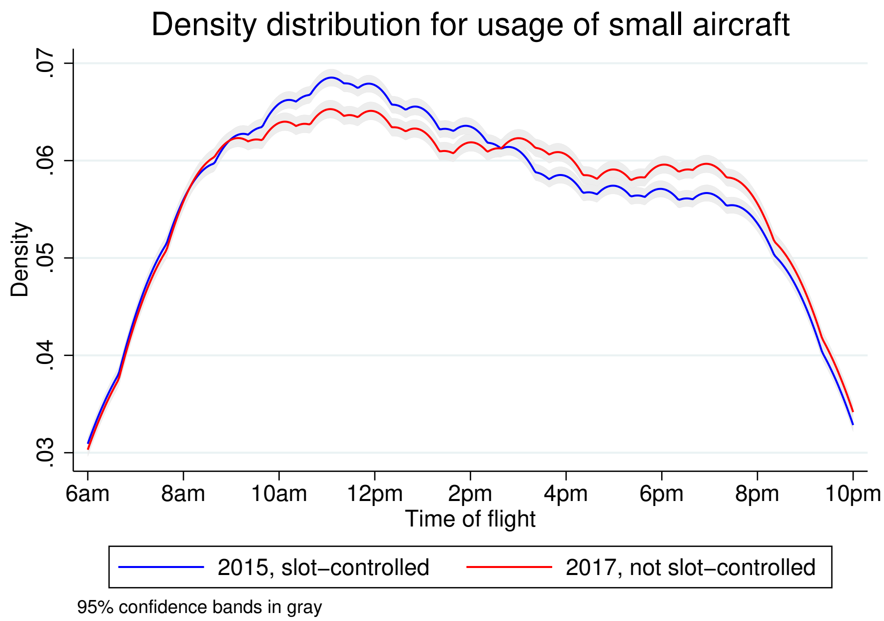

Figure 3: Distribution of small aircraft by time of flight at Newark Airport.

one-hour slot period during the first quarters of 2015 and 2017.[12] The bottom red line is 80% of the top red line and represents the level of usage required to satisfy the use-it-or-lose-it rule.

These graphs should be interpreted cautiously, and together with other suggestive evidence of slot burning, for two reasons. First, we do not have data on each airlines' hourly slot holdings, which may vary significantly from one hour to another.[13] Second, we only observe actual, as opposed to scheduled, departure and arrival times. Both departing and arriving flights experience significant delays, so mapping of a flight into a slot period based on time of departure or arrival could be inaccurate. These factors contribute to occasional non-compliance with the minimum usage requirements or excess of flights over the slot capacity. However, it is clear that airlines under slot controls do not use *all* their slots *all* the time, indicating slack; this slack is not

[12] For United, 869 slots $\cdot$ 90 days/17 hours $= 4,600$ possible flights. For low-cost carriers (Alaska, Jet Blue, Southwest, and Virgin), 92 slots $\cdot$ 90 days/17 hours ≈ 487 possible flights.

[13] We submitted a FOIA request to PANYNJ to get these data.

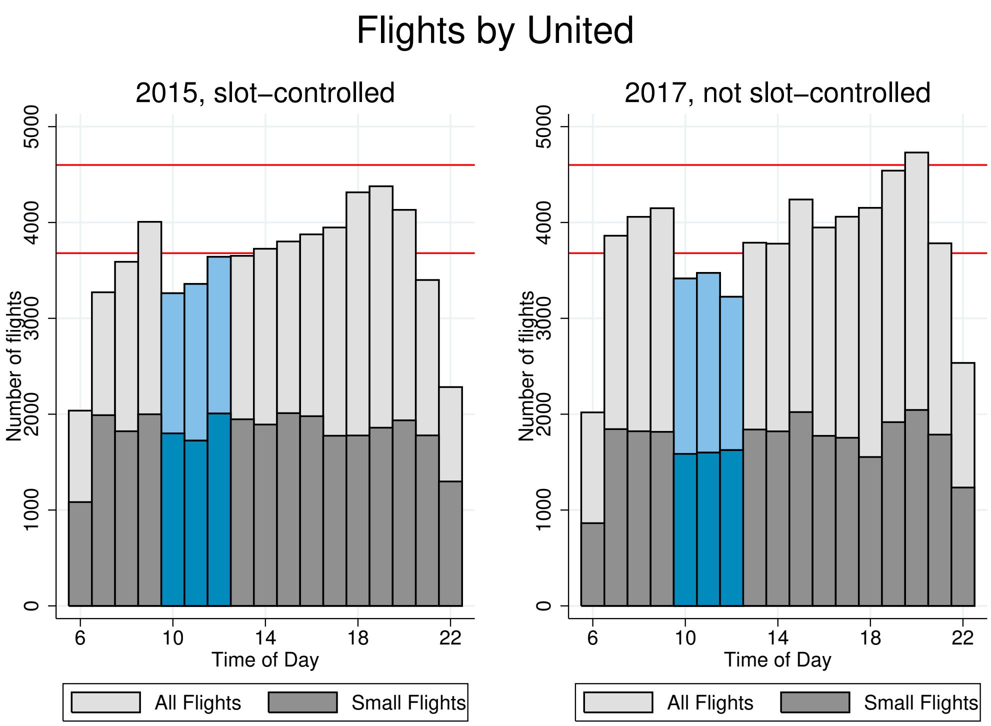

Figure 4: United's compliance with slot and minimum usage constraints.

due to lack of demand, shown by the increase in frequency following the abolition of slot controls. We provide evidence that this trend is explained by slot burning. In particular, we focus on usage of small aircraft during offpeak slot periods (dark blue bars in Figures 4 and 5) that seems to decrease when slot controls are removed in 2017, contrary to the overall increase in flight frequencies.

In section V we refine our descriptive analysis of slot burning, formulate testable hypotheses, and bring them to data. We also investigate the rationing effects of slot control on consumers by looking at the effect of slot restrictions on the number of seats, seat-miles, number of passengers, passenger-miles, and airfare. Since a stated benefit of slot controls is reduced flight delays due to air traffic congestion, we also investigate whether lifting slot restrictions at Newark Airport increases the likelihood of flight delays.

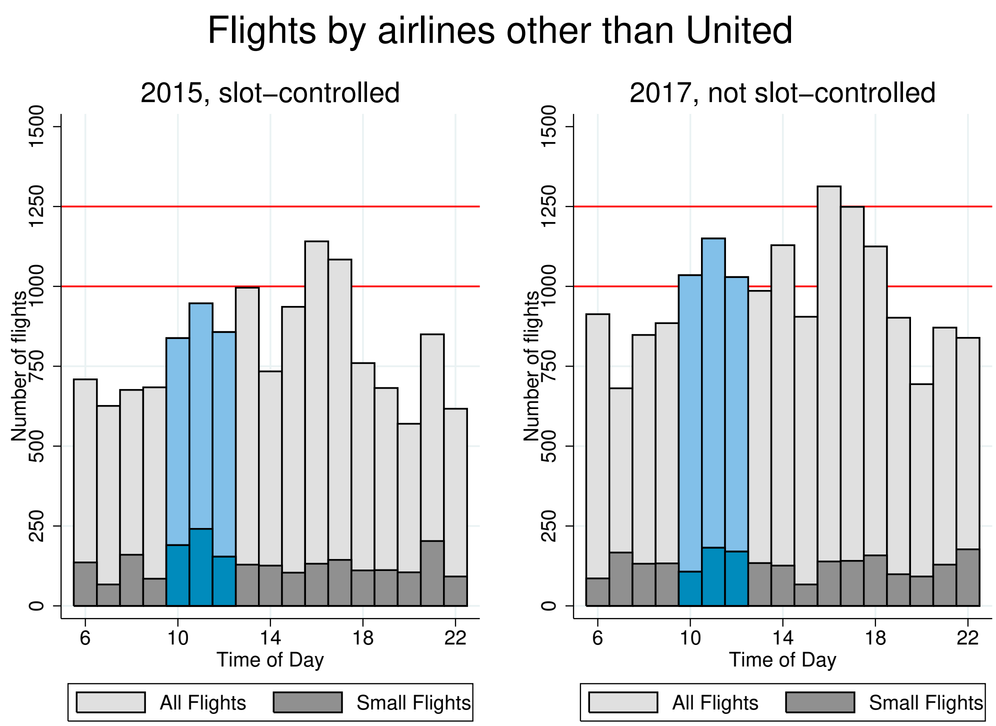

Figure 5: Other airlines' compliance with slot and minimum usage constraints.

III.3 Data and Descriptive Statistics

This paper uses data from three data sources. The main dataset was obtained from the Port Authority of New York and New Jersey (PANYNJ) using FOIA requests and contains information on exact date and time of arrival and departure of all flights to/from Newark Airport, as well as their operating carriers and aircraft types for the 2015-2017 period. We supplement these data with information on the number of passengers and seats for all domestic origin and destination airports at the monthly level from the Air Carrier Statistics (T-100) database by the Bureau of Transportation. Lastly, we use the Airline Origin and Destination Survey (DB1B) from the Bureau of Transportation for the information on ticketing carriers and airfares for a 10% sample of all tickets issued for all domestic itineraries.

Table 2 presents summary statistics of operations at Newark Airport pre- and post-introduction of slot control in 2008 and pre- and post-removal of slot control in 2016. In our empirical analysis, we use the first quarters of each year to control for seasonalities in demand for air travel and scheduling of flights.

Variable	2007		2009		2015		2017	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Small	0.18	0.39	0.26	0.44	0.25	0.43	0.15	0.36
Legacy	0.95	0.21	0.94	0.24	0.81	0.39	0.81	0.39
Delay	35.73	5.16	32.89	5.51	28.19	3.67	31.37	4.54
Load factor	0.78	0.12	0.78	0.11	0.85	0.07	0.87	0.07
Tourist	0.15	0.35	0.14	0.35	0.14	0.36	0.15	0.36
Distance, mi								
< 325	0.20	0.40	0.13	0.33	0.13	0.34	0.12	0.32
325 to 602	0.12	0.32	0.14	0.34	0.11	0.31	0.07	0.26
603 to 998	0.31	0.46	0.30	0.46	0.27	0.44	0.25	0.43
> 998	0.37	0.48	0.43	0.50	0.49	0.50	0.55	0.50
Frequency, daily								
< 6	0.56	0.50	0.81	0.39	0.70	0.46	0.53	0.50
between 6 and 8	0.35	0.48	0.19	0.39	0.22	0.41	0.17	0.38
> 8	0.09	0.29	–	–	0.08	0.28	0.30	0.46
No. routes	24		23		24		24	
No. carriers	17		19		15		16	
No. carrier-routes	59		61		72		75	
No. passengers, mil	1,239.92		1,468.53		1,708.00		2,130.09	
No. flights	17,105,706		12,621,843		12,624,550		16,901,660	

Table 2: Descriptive statistics of operations at Newark Airport in 2007-2017.

As discussed above, slot burning is best evidenced by more frequent usage of small aircraft in relatively less demanded slot periods when an airport is slot-controlled. Unfortunately, we cannot use rigidly defined one-hour slot periods for our analysis to see which slot periods may experience slot burning. Our data provide information on actual arrival and departure times, as opposed to scheduled ones. Given that both departing and arriving flights experience significant delays, the mapping from time of departure or arrival to a slot period is not straightforward. Therefore, in our analysis, we aggregate slot periods into peak and offpeak categories. We define offpeak periods to be between 10am and 1pm, where the demand is sufficiently high to support some flights (unlike the 1am-4am timeframe), but not enough to warrant usage of multiple larger aircraft (unlike a busy timeframe like 7pm-10pm). We refer to any slot-controlled period that is not offpeak as peak. Our results are robust to changing the definition of offpeak periods to 10am-2pm and 11am-2pm.

One could argue that using small aircraft is needed to meet demand for less dense routes (like, for instance, Newark Airport to Martha's Vineyard, MA), and therefore, using small aircraft for these routes should not constitute slot burning. Such justification is harder to accept for dense routes. Our analysis, therefore, is limited to the 28 largest airports in the US by domestic passenger enplanements.[14] Expanding the analysis to all airports (with a dummy variable for the 28 largest airports, which returns a negative coefficient, corroborating the argument described above) yields no meaningful difference in the results to the variables of interest.

We expect airlines that hold large pools of slots at an airport to be more likely to engage in slot burning. We refer to such airlines as slot incumbents. In Newark Airport, slot incumbents are United and regional airlines that operate flights ticketed by United and using United's slots (see Table 1). We incorporate the slot incumbent dummy into our regression analysis in order to control for differential slot-burning incentives of Newark's airlines.

[14] See Appendix A for the list of the top-28 domestic airports.

IV Theoretical Model

In this section we develop a simple theoretical model to inform our empirical analysis and give a formal definition of slot burning.

IV.1 Consumers

For simplicity, assume that consumers purchase only direct flights operated by airline j between a given origin o and destination d airport pair. A consumer i chooses a flight $jodtk$ between airports o and d during a slot period t on an aircraft of capacity k[15] in order to maximize their utility, u_{ijodtk}, given by

$$u_{ijodtk} = x_{jodtk}\beta - \alpha p_{jodtk} + \beta \sum_{k} f_{jodtk} + \xi_{jodtk} + \epsilon_{ijodtk},$$

where x_{jodtk} is a vector of product characteristics, p_{jodtk} is the product airfare, f_{jodtk} is the number of flights available on the service route, ξ_{jodtk} is the unobserved product characteristics, and ϵ_{ijodtk} is the i.i.d. logit error term.

A consumer i chooses product $jodtk$ if their utility from $jodtk$ exceeds the utility from any other product, including the outside good. This results in the standard logit share equations, which we use in the airlines' profit-maximization problems.

[15] We differentiate products by aircraft capacity in order to incorporate capacity choice into the airlines' profit-maximization problems. Consumers may not take the aircraft capacity into account when purchasing flights, so we sum up frequencies of flights between an airport pair over all possible capacities.

IV.2 Airlines

We assume that airlines compete by playing a one-shot game, where they simultaneously choose prices $\mathbf{P}$ and frequencies $\mathbf{F}$ of flights for a given aircraft size k and slot period t[16], subject to the aircraft capacity constraints and slot and use-it-or-lose-it constraints at slot-controlled airports. For simplicity, assume that there are only two aircraft sizes – small and large – so that $k \in \{s, \ell\}$ is a discrete variable that defines capacities of small and large flights. Empirically, we observe that some airlines operate both small and large flights on the same route, e.g. Jet Blue and United, therefore we must allow for this in our theoretical model. To clarify, if an airline does not operate a small flight on route od at slot period t, then $f_{jodts} = 0$.

Let $\boldsymbol{p}_j, \boldsymbol{f}_j$ define the choice variables of airline j for each service route $odtk$ in its set of products $\Omega_j = O_j \times O_j \times T \times K$, where O_j is the set of airports that airline j operates in, T is the set of slot periods, and K is the set of aircraft capacities. Airline j's profit-maximization problem is thus

$$\max_{\boldsymbol{p}_j, \boldsymbol{f}_j} \pi_j \left(\boldsymbol{p}_j, \boldsymbol{f}_j, \boldsymbol{P}_{-j}, \boldsymbol{F}_{-j}\right) = \sum_{odtk} M s_{jodtk} \left(\boldsymbol{P}, \boldsymbol{F}\right) p_{jodtk} - f_{jodtk} MC_{jodtk} - FC_{jodtk}$$

subject to

$$k f_{jodtk} \geq M s_{jodtk} \left(\boldsymbol{P}, \boldsymbol{F}\right) \text{ for all } odtk \in \Omega_j \tag{1}$$

$$\sum_{dk} f_{jodtk} + \sum_{dk} f_{jdotk} \leq S_{jot} \text{ for all } ot \in O \times T \tag{2}$$

$$\sum_{dk} f_{jodtk} + \sum_{dk} f_{jdotk} \geq 0.8 S_{jot} \text{ for all } ot \in O \times T \tag{3}$$

$$f_{jodtk} \geq 0 \text{ for all } odtk \in \Omega_j \tag{4}$$

[16] As a reminder, a slot period at Newark Airport is defined to be a 30-minute window on a particular day of week during a scheduling season, e.g. 6:00-6:30 am on Mondays during the winter season of 2015.

Inequality (1) describes the aircraft capacity constraint a for each product $odtk$ in Ω_j[17], (2) is the slot constraint with $S_{jot} = \infty$ for $ot \in O \times T$ if airport o is not slot-controlled in period t, (3) is the use-it-or-lose-it (UIOLI) constraint, and (4) is the non-negativity constraint on the frequency of flights.

The Lagrangian is then

$$\mathcal{L} = \sum_{odtk} Ms_{jodtk}\left(\boldsymbol{P},\boldsymbol{F}\right)p_{jodtk} - f_{jodtk}MC_{jodtk} - FC_{jodtk} - \gamma_{jodtk}(Ms_{jodtk}(\boldsymbol{P},\boldsymbol{F}) - kf_{jodtk}) +$$

$$+\lambda^s_{jot}\left(S_{jot} - \sum_{dk} f_{jodtk} + \sum_{dk} f_{jdotk}\right) - \lambda^u_{jot}\left(0.8S_{jot} - \sum_{dk} f_{jodtk} + \sum_{dk} f_{jdotk}\right)$$

and the FOCs simplify to

$$s_{jodtk} + \sum_{\widetilde{odtk}\in\Omega_j} \frac{\partial s_{j\widetilde{odtk}}}{\partial p_{jodtk}}\left(p_{j\widetilde{odtk}} - \gamma_{j\widetilde{odtk}}\right) = 0$$

$$M\sum_{\widetilde{odtk}\in\Omega_j} \frac{\partial s_{j\widetilde{odtk}}}{\partial f_{jodtk}}\left(p_{j\widetilde{odtk}} - \gamma_{j\widetilde{odtk}}\right) - MC_{jodtk} + \gamma_{jodtk}k + \lambda^u_{jot} - \lambda^s_{jot} = 0$$

We distinguish two types of slot periods – peak and offpeak periods. Peak periods are periods during which the slot constraint is binding, so $\lambda^s_{jot} > 0$, and the UIOLI constraint is automatically satisfied, so $\lambda^u_{jot} = 0$. In other words, airlines do not need to burn slots in order to satisfy the minimum usage requirements during the peak periods. In contrast, offpeak periods are periods during which the slot constraint is slack, so $\lambda^s_{jot} = 0$, and the UIOLI constraint is binding, so $\lambda^u_{jot} > 0$. Therefore, in order to satisfy the minimum usage requirements during offpeak periods, airlines

[17] We abstract away from potential fleet constraints. It is not clear how to model them since legacy airlines are known to hire regional airlines to operate flights ticketed through their booking systems.

must burn slots.

Definition 1. Airline j burns slots at airport o during a slot period t if the slot constraint is slack and the UIOLI constraint binds, or $\lambda^s_{jot} = 0$ and $\lambda^u_{jot} > 0$.

Our theoretical model predicts different responses of flight frequencies in peak- and offpeak slot periods after a removal of slot control. In conjunction with Figures 4 and 5, that establish airlines and slot periods with binding slot and UIOLI constraints, we are able to draw corollaries to test for evidence of slot burning using the reduced-form approach.[18]

V Empirical Analysis and Results

This section introduces testable hypotheses from the theoretical model with the corresponding regression specifications and discusses the results.

V.1 Frequency

Consider a counterfactual where slot control is removed, i.e. $\lambda^s_{jot} = 0$ and $\lambda^u_{jot} = 0$. We expect flight frequencies to increase during the peak slot periods and decrease during the offpeak periods, holding all else equal.[19] We also expect the effect to be more pronounced for legacy airlines because they hold more slots than low-cost carriers. We explore this hypothesis in regression specification (1).

Hypothesis 1. *After slot control removal, the frequency of flights increases in the peak and decreases in the offpeak slot periods, more so for slot incumbents, holding*

[18] We do not attempt to recover the values of λ^s and λ^u. This endeavor is left for future work.

[19] Removal of slot control rules decreases entry costs for new airlines, in particular, low-cost carriers. Changes in frequencies and composition of flights in response to entry post slot control are a part of the changes due to removal of the slot and UIOLI constraints per se, since removal of slot control eliminates the foreclosure incentive that causes slot burning in the first place.

all else fixed.

Model 1. The model tests the change in flight frequency due to slot-restrictions,, including any differential changes between peak and offpeak hours, by using flights between Newark Airport and the 28 largest airports in the US for the first quarters of 2015 and 2017. The years were chosen to fall on two sides of the 2016 reclassification of Newark Airport from a slot-controlled (Level 3) airport to a schedule-facilitated (Level 2) airport. The proprietary data comes from FOIA requests to the Port Authority of New York and New Jersey (PANYNJ) to discern whether a flight takes place during peak hours or offpeak hours.

$$
\begin{aligned}
freq(\text{-miles}) = &\beta_1 slot + \beta_2 incumbent + \beta_3 slot \times incumbent + \beta_4 \mathit{offpeak} + \beta_5 \mathit{offpeak} \times incumbent \\
&+ \beta_6 \mathit{offpeak} \times slot + \beta_7 \mathit{offpeak} \times incumbent \times slot + \beta_i (controls)
\end{aligned}
$$

We include fixed effects for distance (binned at less than 325 miles, between 325 and 602 miles, between 603 and 998 miles, and more than 998 miles), following GAO (2012) specifications, and airport fixed-effects for the non-Newark airport in the origin-destination pair. We use data from the first quarters of 2015 and 2017 to control for seasonality.

Discussion, Columns 1. First, we compare the number of flights before and after slot-controls by carrier incumbency and time of day (peak or offpeak). Following the removal of slot-controls, the incumbent increases the number of flights during peak hours ($p=0.0000$), but operates roughly the same number of flights during offpeak hours ($p=0.2102$). Thus, our prediction for peak hours is validated, but cannot be confirmed for offpeak hours. Controlling for distance and airports, non-incumbents offer more flights during both peak and offpeak hours when Newark Airport is slot-controlled.

Next, we compare the number of flights during peak and offpeak periods by carrier incumbency and slot regime. We observe that non-slot-incumbents fly more flights during peak hours than offpeak hours when Newark Airport is not slot-controlled (*p=0.0000*). However, they fly roughly the same number of flights (*p=0.4513*) when Newark Airport becomes slot-controlled. This conduct can be evidence of slot burning, or explained by the fact that the slot constraint binds during the peak hours, forcing the airlines to reallocate their flights to offpeak hours.

We also observe slot incumbents have more flights during peak hours than offpeak hours when Newark Airport is not slot-controlled (*p=0.0255*). However, the incumbents fly *more* flights during offpeak hours than during peak hours when Newark Airport becomes slot-controlled (*p=0.0030*). A binding slot constraint during peak hours cannot explain why they operate *more* flights during offpeak hours when Newark Airport is slot-constrained. This conduct is indicative of slot burning.

While the number of flights operated by the incumbent during offpeak period is roughly the same regardless of whether Newark Airport is slot-controlled or not (*p=0.2102*), there is a decline in *frequency-miles* when Newark becomes slot-controlled (Column 1b). This confirms our hypothesis that slot-burning happens along shorter routes.

Lastly, we note that frequency of flights is uncorrelated with measures of quantity (either seats or passengers), which are discussed in Subsection V.5.

V.2 Use of Small Flights

Slot burning implies that airlines are operating loss-making flights in the offpeak periods. The best approach to minimizing said loss is to (i) fly a smaller aircraft, (ii) across a shorter distance, (iii) on a route with relatively high demand. In our regression analysis, we restrict our sample to the 28 largest airports by passenger enplanements in the contiguous United States and control for the flight distance in order to hold factors (ii) and (iii) fixed. Therefore, usage of small aircraft at slot- and non-slot-controlled airports can be used as a proxy for slot burning. We explore this hypothesis in regression specification (2). As with Hypothesis 1, we expect the effect to be stronger for slot incumbents.

Hypothesis 2. *Under slot control, usage of small aircraft is more prevalent during the offpeak slot periods, more so for slot incumbents, holding all else fixed.*

Model 2. The model tests if the probability of using small aircraft during peak and offpeak periods for the slot incumbent and non-slot incumbents changes when Newark's slot control is removed. The independent variable is whether Newark Airport is slot-controlled (2015)[20] or not (2017).

$$\begin{aligned}\mathbb{L}_{small} =& \beta_1 slot + \beta_2 incumbent + \beta_3 slot \times incumbent + \beta_4 offpeak + \beta_5 offpeak \times incumbent \\ &+ \beta_6 offpeak \times slot + \beta_7 offpeak \times incumbent \times slot + \beta_i (controls)\end{aligned}$$

The dependent variable is an indicator for small aircraft, defined as aircraft carrying 100 passengers or less. This definition is to be consistent with the GAO's

[20]Only flights between 6am and 10:59am are slot-controlled at Newark Airport, and are coded as such in the data.

(2012) model; we try different definitions of "small" (for instance, aircraft carrying fewer than 81 passengers), to no meaningful change in the coefficients. The $\mathbb{L}$ denotes log-odds.

We include the distance and airport fixed effects as before. Additionally, we use fixed effects for daily flight frequency (binned at less than 6, between 6 and 8, and more than 8 flights), following GAO (2012) specifications.

As a robustness check, we also add data from 2007 and 2009, to encompass before and after Newark Airport transitioned to being a slot-controlled airport. We find that the qualitative results remain the same – that all carriers are more likely to use small aircraft in slot-controlled airports. However, given the severe change in air travel demand following the financial crisis, we decided against including the years 2007-2009 into our main (or any other) specification.

Discussion, Column 2. We find that slot incumbents, under no restrictions, were 20% more likely to use small aircrafts during offpeak hours (than during peak hours), a statistic that jumps to 40% under slot restrictions ($p = 0.0213$).[21] That the slot incumbent is twice as likely to use small aircrafts during offpeak hours (than peak hours) when Newark Airport becomes slot-controlled is indicative of slot burning. Similarly, non-incumbents exhibit an increased reliance on small aircrafts when Newark Airport is slot-controlled ($p = 0.0012$). The difference between the two odds ratios for all carriers shows that the increased usage of small aircraft during offpeak hours cannot only be explained by the type of consumer demand during

[21] From Model 2, the likelihood of slot incumbent to use small aircraft during offpeak hours when Newark Airport is slot-controlled = (Coefficient on $slot_1 \times incumbent_1 \times offpeak_1$)/(Coefficient on $slot_1 \times incumbent_1 \times offpeak_0$) = 7.772/ 5.527 = 1.406. The same exercise for when incumbents fly without slot restrictions yields 1.216.

offpeak hours. Figure 3 in Section III illustrates that usage of small flights increases during offpeak hours (for example, 10am-1pm) and decreases during peak hours (for example, 7pm to 10pm) when Newark Airport is slot-controlled.

V.3 Consecutive Flights

A likely argument in defense of airlines' frequent use of small aircraft is catering to the time-sensitivity of demand. We address this concern by looking at the probability of observing consecutive flights to the same destination within a short timeframe (30 minutes in the baseline specification).

Hypothesis 3. *Under slot control, the probability of observing consecutive flights to the same destination is greater in the offpeak slot periods, more so for slot incumbents, holding all else fixed.*

We explore this hypothesis in regression specification (3a). In regression specification (3b), we further refine our test by examining the probability of observing consecutive *small* flights to the same destination within a short timeframe.

Model 3a. The second model investigates the probability of an airline flying multiple flights on a route within a 30-minute window. The dependent variable indicates whether the same airline offers another flight from the same origin to the same destination within the 30-minute window. The qualitative results are the same if the window is changed to 45- or 60-minutes. The same set of fixed effects are used from the previous model.

$$\mathbb{L}_{consecutive} = \beta_1 slot + \beta_2 incumbent + \beta_3 slot \times incumbent + \beta_4 offpeak + \beta_5 offpeak \times incumbent + \beta_6 offpeak \times slot + \beta_7 offpeak \times incumbent \times slot + \beta_i(controls)$$

Model 3b. This model investigates the probability of an airline flying multiple *small* flights on a route within a 30-minute window. This specification helps us narrow down the mechanism by which airlines burn slots (that is, whether airlines are indeed flying multiple small flights during offpeak hours).

$$\mathbb{L}_{consec\_small} = \beta_1 slot + \beta_2 legacy + \beta_3 slot \times legacy + \beta_4 offpeak + \beta_5 offpeak \times legacy + \beta_6 offpeak \times slot + \beta_7 offpeak \times legacy \times slot + \beta_i(controls)$$

Discussion, Columns 3. We find that while non-incumbents do not significantly crowd their flights under slot restrictions, slot incumbents are 11% more likely to do so when operating out of Newark Airport under slot controls than without slot controls ($p = 0.0429$). The definition of a consecutive flight in the main specification is another flight within a 30-minute window; the qualitative results are robust to alternate definitions of consecutive (45- or 60-minutes).

We hypothesized that refining our test by examining if the probability of observing consecutive *small* flights to the same destination will yield similar results. In fact, slot incumbents are 75% more likely to fly consecutive small flights when Newark Airport is slot-controlled ($p=0.0000$); non-incumbent carriers are almost twice as likely to fly consecutive small flights when Newark Airport is slot-controlled ($p=0.0110$). These results suggest that all carriers increase their reliance on small consecutive flights when Newark Airport becomes slot-controlled, which is indicative of slot burning.

V.4 Delay

Although not from the theoretical model, the rationale provided by the FAA for slot restrictions invokes congestion, leading us to believe that slot controls may alleviate delays.

Hypothesis 4. *Slot control improves airlines' on-time performance.*

Model 4. This model investigates the probability that a flight is delayed by longer than 30 minutes. The data comes from the T-100 database, which contains aggregate information on flight schedules. We use airport-specific dummies to account for the possibility that delays could be caused by congestion at the other endpoint airport.

$$\mathbb{L}_{delay} = \beta_1 slot + \beta_2 incumbent + \beta_3 slot \times incumbent + \beta_i(controls)$$

Discussion, Column 4. The relegation of Newark Airport to a Level 2 airport resulted in a worse on-time performance for Newark Airport. Column 4 shows that while slot incumbents fare slightly worse in terms of on-time performances in 2017 (when Newark Airport was not slot-controlled), all airlines experience less delays in 2015. The probability of a legacy carrier being delayed by 30 minutes or more decreases by about 17%,[22] while non-incumbent carriers are 68% less likely to be delayed in 2015.

[22] 1 - 0.850/1.026 = 0.172.

	(1a)	(1b)	(2)	(3a)	(3b)
	Frequency	Freq-mile	Small	Consecutive	SmallConsec
Model	OLS	OLS	Logistic	Logistic	Logistic
$slot_0 \times incumbent_0 \times offpeak_1$	-0.828***	-245.4*	0.874	0.654**	0.265
	(-9.99)	(-1.97)	(-1.70)	(-3.14)	(-1.79)
$slot_0 \times incumbent_1 \times offpeak_0$	5.577***	8554.0***	9.979***	1.358***	3.192***
	(126.28)	(128.61)	(44.99)	(4.57)	(4.58)
$slot_0 \times incumbent_1 \times offpeak_1$	5.424***	7620.1***	12.14***	0.754**	1.166
	(74.36)	(69.37)	(39.60)	(-3.14)	(0.52)
$slot_1 \times incumbent_0 \times offpeak_0$	1.206***	778.2***	1.392***	1.048	0.782
	(22.71)	(9.73)	(5.81)	(0.63)	(-0.76)
$slot_1 \times incumbent_0 \times offpeak_1$	1.272***	1435.3***	1.804***	0.704*	1.681
	(14.31)	(10.72)	(6.70)	(-2.49)	(1.48)
$slot_1 \times incumbent_1 \times offpeak_0$	5.114***	6672.3***	5.527***	1.430***	2.967***
	(111.81)	(96.88)	(34.24)	(5.28)	(4.27)
$slot_1 \times incumbent_1 \times offpeak_1$	5.310***	6636.2***	7.772***	0.955	3.251***
	(71.71)	(59.50)	(33.37)	(-0.54)	(4.41)
Controls:					
Distance	Yes	Yes	Yes	Yes	Yes
Airport	Yes	Yes	Yes	Yes	Yes
Frequency			Yes	Yes	Yes
N	77,776	77,776	77,776	77,776	21,458

t statistics in parentheses

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

Table 3: Regression coefficients

	(4) Delayed	(5a) Seats (thousands)	(5b) Seat-mile (millions)	(6a) Passengers (thousands)	(6b) Pax-mile (millions)	(7) Mktfare ($)
Model	Logistic	OLS	OLS	OLS	OLS	OLS
$slot_0 \times incumbent_1$	1.026	-66.42***	-54.56***	-52.83***	-44.39***	41.00***
	(0.90)	(-22.16)	(-12.80)	(-21.04)	(-12.16)	(79.87)
$slot_1 \times incumbent_0$	0.321***	-12.90	-9.028	-9.293	-8.504	13.70***
	(-32.49)	(-1.73)	(-0.85)	(-1.48)	(-0.93)	(16.98)
$slot_1 \times incumbent_1$	0.850***	-50.82 ***	-36.31***	-39.74***	-30.74**	46.58***
	(-6.02)	(-6.60)	(-3.32)	(-6.16)	(-3.28)	(55.97)
Constant		231.0***	89.69***	194.7***	76.55***	291.4***
		(14.55)	(3.97)	(14.64)	(3.96)	(117.30)
Controls:						
Distance	Yes	Yes	Yes	Yes	Yes	Yes
Airport	Yes	Yes	Yes	Yes	Yes	Yes
Frequency	Yes	Yes	Yes	Yes	Yes	Yes
Region	Yes					Yes
N	71,099	1,989	1,989	1,989	1,989	1,419,871

t statistics in parentheses

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

Table 4: Regression coefficients (cont'd)

V.5 Quantity and Price

We have no testable hypotheses for quantity or price. Changes in price, capacity, and quantity of passengers are all ambiguous and dependent on various intertwined factors. For illustration, assume that under no slot restrictions, an airline flies one 150-passenger flight along a route. Following slot restrictions, the airline might choose to fly two or three 60-passenger flights, depending on load factors and passenger sensitivity to frequency. Without knowing the direction of the change in supply, it is not possible to know, a priori, the direction of the change in price. Questions relating to consumer surplus, therefore, can only be answered by the data.

Models 5-6. The subsequent models estimate the impact of slot restrictions on the number of seats available (Model 5) and the number of passengers transported (6). The T-100 database is used for these estimations, since it comes with aggregate numbers for seats and passengers. Since T-100 database does not include time of flight, thus whether a flight is traveling during peak or offpeak hours cannot be included in the model. Controls include fixed effects for frequency, distance, year, region, and airport.

$$seat(\text{-miles}) \text{ or } passenger(\text{-miles}) = \beta_0+\beta_1 slot+\beta_2 incumbent+\beta_3 slot\times incumbent+\beta_i(controls)$$

Model 7. The seventh model looks at the effect on price at slot-controlled airports. We use any routes within the top-28 airports for this analysis, with the *slot* dummy indicating whether any of the airports within the route were slot-controlled in the period the flights took place. By restricting the analysis to the first quarters of 2015 and 2017, we can exploit the reclassification of Newark Airport as an exogenous variation of the independent variable, *slot*. Since we use the DB1B database for this analysis, we do not compute any measures for frequency of flights or time of flight within the route. Controls include fixed effects for distance, year, region, and airport.

$$price = \beta_0 + \beta_1 slot + \beta_2 incumbent + \beta_3 slot \times incumbent + \beta_i(controls)$$

Discussion, Columns 5-7. We find an increase in the total number of seats flown by the slot incumbent (Column 5a, $p= 0.0426$) in presence of slot control, but we cannot reject the null hypothesis that the incumbent flies different seat-miles at the 5% level. This shows that while the incumbent is offering more seats, the seats are on shorter routes, which alludes to slot-burning. Non-incumbents do not show any statistically significant change in seats or seat-miles.

Similarly, incumbents fly more passengers (Column 6a, p=*0.0422*) following slot restrictions, but exhibit no change in passenger-miles. Non-incumbents fly the same number of passengers and passenger-miles.

Column 7 outlines the effect of slot restrictions on price. We find that the slot incumbent (non-incumbents) charge about $5 more ($13 more) per passenger for one-way travel when they serve a slot-controlled airport, with the median ticket price being around $220. These two values are statistically significant (p=*0.0000* for both incumbent and non-incumbents), and different from one another (p= *0.0000*).

V.6 Discussion on Consumer Surplus

The effect of slot restrictions on consumer surplus is ambiguous. Holding all else equal, more passengers fly under slot control but on average pay a higher price. Such an effect is indicative of an outward shift in the demand curve, likely suggesting that the product quality has increased. We also find that the total number of seats offered under slot control increases, suggesting that there is an outward shift in the supply curve as well. The fact that we observe an increase in the number of passengers transported together with an increase in airfares implies that the shift in the demand curve is larger than that in the supply curve.

We attribute the increase in product quality to the success of slot control at curbing delays. Ceteris paribus, a flight with poor on-time performance is considered inferior to one that reliably follows schedule. However, as we document in the next section, removing slot control does add new direct routes for certain communities, decreasing layover times and acting as a counterveiling force to the increase in product quality associated with slot control. We rationalize this by the fact that a consumer's disutility from delay may be different from their disutility from a long layover or poor match between their desired and actual departure time, since the delay is an uncertainty only realized at the point of use, not the point of purchase. Moreover, frequency of flights on a route is another determinant of product quality. All else constant, a route with frequent flights provides a better match between a time-sensitive consumer's desired departure time and the actual departure time. A structural model is needed to quantify the relative tradeoffs and we leave this for future work.

Not only is the effect on consumer surplus ambiguous, it is also heterogeneous. Our findings show that slot restrictions increase the number of flights along dense business-routes in short distances (the routes that facilitate slot burning), while decreasing the number of flights to tourist destinations. Therefore, we expect differential impact on consumers flying different routes; consumers flying tourist routes experience a reduction in consumer surplus (higher price, smaller number of seats available, lower quality product), whereas the effect on business destinations are ambiguous (higher price indicates lower surplus, but a higher product quality indicates higher surplus).

VI Entry and Exit Following Reclassification

In this section, we analyze airlines' entry and exit decisions at the route-level following the Newark's removal of slot controls. We also study entry and exit in 2019 in order to benchmark the magnitude of the industry shake-out in 2016.[23] As before, we limit our

[23]We use the T-100 database from the Bureau of Transportation Statistics. As of now, January 2020 is the latest available data, but even once more data become available, we could not use 2020 due to the impacts of the COVID-19 pandemic on air travel.

analysis to the first quarters of 2015 and 2017-2019 to overstep potential seasonality issues. However, given our interest in understanding the differential impact of slot control on communities of varying size, and especially in small and rural communities' access to air transport, we extend our sample to all domestic airports of the contiguous United States.

	Number of routes				2017		2019	
	2015	2017	2018	2019	Entry	Exit	Entry	Exit
American	5	6	5	5	1	0	0	0
Allegiant	0	4	1	3	4	0	2	0
Alaska	1	4	4	6	3	0	2	0
Delta	5	5	5	5	0	0	0	0
JetBlue	6	6	7	7	0	0	0	0
Elite	0	1	1	0	1	0	0	1
Spirit	0	3	6	8	3	0	2	0
Southwest	8	7	8	9	2	3	3	2
United	77	83	83	77	9	3	4	10
Virgin	2	2	2	0	0	0	0	2
Total	104	121	122	120	23	6	13	15

Table 5: Summary of airlines' entry and exit decisions in 2015-2017 and 2018-2019.

Table 5 above summarizes airlines' entry and exit decisions. Only seven airlines were operating in Newark in 2015. The reclassification brought in three new airlines: Allegiant Air (a low-cost carrier based in Nevada), Elite Airways (a brand new airline operating out of Portland, ME), and Spirit. Not all airlines entered new markets; Delta, JetBlue, and Virgin did not change their routes at all. Of the seven existing airlines only two, Southwest and United, dropped routes. Overall, Newark Airport saw 23 entry events on 19 routes and six exits on six routes in 2017. Between 2018 and 2019, the latest available years with no change in slot regime and unrestricted entry, Newark experienced significantly less entry (13 entry events on 13 routes) and more exit (15 exits on 15 routes).

	Alaska	Allegiant	American	Elite	Southwest	Spirit	United	No. carriers, 2017
Akron, OH	–	–	–	–	–	–	1	1
Alcoa, TN	–	1	–	–	–	–	–	2
Asheville, NC	–	1	–	–	–	–	1	2
Binghamton, NY	–	–	–	–	–	–	-1	0
Chattanooga, TN	–	–	–	–	–	–	1	1
Chicago, IL	–	–	1	–	–	–	–	2
Flint, MI	–	–	–	–	–	–	1	1
Fort Drum, NY[24]	–	–	–	–	–	–	-1	1
Fort Lauderdale, FL	–	–	–	–	1	1	–	4
Fort Wayne, IN	–	–	–	–	–	–	1	1
Hebron, KY	–	1	–	–	–	–	–	3
Houston, TX	–	–	–	–	-1	–	–	0
Kenner, LA	–	–	–	–	-1	–	–	1
Key West, FL	–	–	–	–	–	–	1	1
Lake City, FL	–	–	–	–	–	–	-1	0
Lexington, KY	–	–	–	–	–	–	1	1
Myrtle Beach, SC	–	–	–	–	–	1	–	2
Nashville, TN	–	–	–	–	-1	–	–	1
Orlando, FL	–	–	–	–	1	1	–	4
Portland, OR	1	–	–	–	–	–	–	2
Salt Lake City, UT	–	–	–	–	–	–	1	2
San Diego, CA	1	–	–	–	–	–	–	2
San Jose, CA	1	–	–	–	–	–	1	2
Savannah, GA	–	1	–	–	–	–	–	2
Vero Beach, FL	–	–	–	1	–	–	–	1
No. mkts, 2015	1	0	5	0	8	0	77	
No. mkts entered—exited	3—0	4—0	1—0	1—0	2—3	3—0	9—3	

Table 6: All entry and exit decisions by airline between 2015 and 2017.

Tables 6 and 7 below provide detailed information on routes that experienced entry (encoded as 1) and exit (-1) following the reclassification in 2016, and in 2019. Removal of slot control resulted in entry on 19 routes, seven of which are brand new (highlighted in bold in Table 6), with six of them due to United. Interestingly, low-cost carriers did not start new routes; instead, they entered mid-sized (Alcoa, TN; Asheville, NC; Hebron, KY; Savannah, GA; Myrtle Beach, SC) and West Coast (Portland, OR; San Diego, CA; San Jose, CA) airports, challenging United on those routes. Another group of airports that experienced entry is tourist destinations in Florida. In this case, an entry by a low-cost carrier was matched by another low-cost carrier and challenged United and Jet Blue, resulting in a four-firm oligopoly (Fort Lauderdale, FL and Orlando, FL).

Airlines exited from six markets in 2017, stopping operations on three routes (Binghamton, NY; Houston, TX; Lake City, FL). Out of all exits, only one occurred on a route from the slot-burning sample – Houston, TX by Southwest. Therefore, we can conclude that the slot incumbent did not operate entirely unprofitable routes just for the sake of burning slots. Moreover, Table 6 is suggestive for refuting anecdotal claims that the FAA may be tolerating slot burning if slots are burned on routes providing access to air transport for small and rural communities. Slot controls alone do not appear to create incentives for airlines to operate flights to small airports as evidenced by lack of mass exit from small destinations. Generally, the fact that the reclassification resulted in entry into a variety of destinations, and exit from a handful of destinations, implies that a heterogeneous group of consumers have benefited from the change in slot regime.

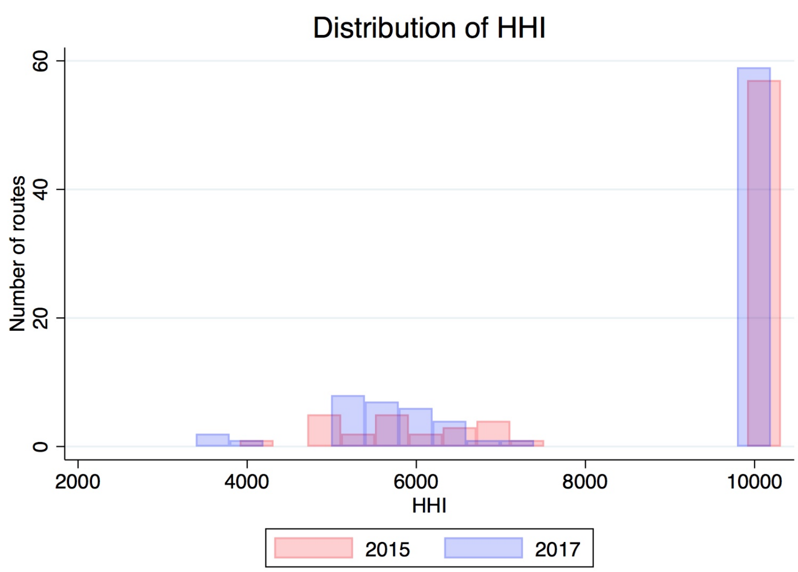

Figure 6: Distribution of HHI on routes from/to Newark in 2015 and 2017.

All in all, changes in frequencies of flights on the extensive and the intensive margin contributed to a slight leftward shift of the HHI distribution. See Figure 6. The average HHI in 2015 is around 8,211 relative to 8,427 in 2017. However, this seeming increase in concentration is due to United opening six new monopoly routes. Conditional on existence of routes in 2015, the average HHI declines from 7,903 in 2015 to 5,291 in 2017.

	Alaska	Allegiant	Elite	Southwest	Spirit	United	Virgin	No. carriers, 2019
Alcoa, TN	–	1	–	–	–	–	–	2
Asheville, NC	–	1	–	–	–	–	–	2
Atlanta, GA	–	–	–	–	1	–	–	3
Avoca, PA	–	–	–	–	–	-1	–	0
Baltimore, MD	–	–	–	–	–	-1	–	0
Chattanooga, TN	–	–	–	–	–	-1	–	0
Des Moines, IA	–	–	–	–	–	-1	–	0
Fort Wayne, IN	–	–	–	–	–	-1	–	0
Horseheads, NY	–	–	–	–	–	1	–	1
Indianapolis, IN	–	–	–	-1	–	–	–	1
Ithaca, NY	–	–	–	–	–	-1	–	0
Los Angeles, CA	1	–	–	–	–	–	-1	2
Montrose, CO	–	–	–	–	–	1	–	1
Myrtle Beach, SC	–	–	–	–	–	-1	–	1
Nashville, TN	–	–	–	1	–	–	–	2
Oakland, CA	–	–	–	1	–	–	–	1
Orlando, FL	–	–	–	-1	–	–	–	3
Palm Springs, CA	–	–	–	–	–	1	–	1
Presque Isle, ME	–	–	–	–	–	1	–	1
San Diego, CA	–	–	–	1	–	–	–	3
San Francisco, CA	1	–	–	–	–	–	-1	2
San Jose, CA	–	–	–	–	–	-1	–	1
South Bend, IN	–	–	–	–	–	-1	–	0
Tampa, FL	–	–	–	–	1	–	–	3
Vero Beach, FL	–	–	-1	–	–	–	–	0
Windsor Locks, CT	–	–	–	–	–	-1	–	0
No. mkts, 2018	4	1	1	8	6	83	2	
No. mkts entered—exited	2—0	2—0	0—1	3—2	2—0	4—10	0—2	

Table 7: All entry and exit decisions by airline between 2018 and 2019.

Tracking entry patterns over time, we document that a recent entrant, Elite Airways, ceased operations in Newark by 2019, so did Virgin by dropping the Los Angeles and San Francisco routes that were entered by Alaska the same year. United's exit from eight markets (Baltimore, MD; Chattanooga, TN; Des Moines, IA; Fort

Wayne, IN; Ithaca, NY; South Bend, IN; Windsor Locks, CT) completely stopped operations on these routes; two of those routes (Chattanooga, TN and Fort Wayne, IN) were entered in 2017, after removal of slot control. Additionally, United exited another two routes that experienced entry in 2017 – Myrtle Beach, SC by Spirit and San Jose, CA by Alaska. The fact that many of the routes with entry in 2017 experienced exit in 2019 implies that the industry was still underway to the long-run equilibrium. The average HHI in 2018 and 2019 were 7,398 and 7,419, respectively.

It is possible that the exits in 2019 are delayed decisions due to the change in slot rules in 2016. If it were the case, then the shutdown of some routes altogether is a concerning effect of the slot liberalization. However, the number of routes shut down is still small, and all the towns losing direct service to/from Newark Airport have (i) connecting flights to Newark Airport, and (ii) are within 100 miles of another airport with direct service to Newark Airport. All of this suggests that even if the 2019 exits were due to the 2016 change in slot rules, its effects are minimal.

Comparing the 2015 and 2019 figures, we can conclude that removal of slot control at Newark brought in a competitive low-cost carrier, Spirit, and 16 additional carrier-routes thereby decreasing concentration by nearly 10%. Thus, we can conclude that removal of slot control was favorable to promoting competition at the Newark Airport.

VII Conclusion

In this study, we show that firms respond to slot restrictions by using smaller flights to use their allocated slots in order to meet the usage requirements. Eliminating such restrictions results in entry, primarily from newly formed low-cost carriers with no historic footprint at the airport. This entry can lower prices, but also result in flight delays due to congestion. However, more passengers fly when Newark Airport is slot-controlled (along routes used to burn slots), implying an ambiguous change in consumer surplus following slot liberalization (lower price, fewer passengers, more delays).

However, since low-cost entrants offer a different assortment of products than the incumbent (almost always a legacy carrier), any change in the relative balance between the two will have differential impact on passengers flying routes dominated by legacy or low-cost carriers. Due to this heterogenous effect on conusmers, policy decisions on slot restrictions to manage congestion at airports must be balanced with an eye on the foreclosure incentive by airlines, and subsequent changes in consumer welfare due to changes in product quality (frequency on a route) and the price paid by passengers.

References

Ali, R., 2020. Effect of Airport Slots in Competition and Antitrust Policy: Evidence from a Recent Merger. *Working paper.*

Ciliberto, F. and Williams, J.W., 2010. Limited access to airport facilities and market power in the airline industry. *The Journal of Law and Economics*, 53(3), pp.467-495.

Fukui, H., 2019. How do slot restrictions affect airfares? New evidence from the US airline industry. *Economics of Transportation*, 17, pp.51-71.

Government Accountability Office (GAO), 2012. Slot-Controlled Airports: FAA's Rules Could be Improved to Enhance Competition and Use of Available Capacity. Report GAO-12-902.

Luttman, A. 2019. Capacity constraints and service quality: Do airport slot controls reduce flight delays? Accessed at https://sites.google.com/a/uci.edu/aluttmann/research?authuser=0.

Pallini, T. “FAA moves to prevent ’ghost flights’ of near-empty planes in the US by easing rules that put them at risk of losing their spots at airports.” *Business Insider*, 11 March 2020, https://www.businessinsider.com/coronavirus-faa-suspending-slot-requirements-at-us-airports-2020-3.

Peterson, E.B., Neels, K., Barczi, N. and Graham, T., 2013. The economic cost of airline flight delay. *Journal of Transport Economics and Policy (JTEP)*, 47(1), pp.107-121.

Snider, C. and Williams, J.W., 2015. Barriers to entry in the airline industry: A multidimensional regression-discontinuity analysis of AIR-21. *Review of Economics and Statistics*, 97(5), pp.1002-1022.

Swaroop, P., Zou, B., Ball, M.O. and Hansen, M., 2012. Do more US airports need slot controls? A welfare based approach to determine slot levels. *Transportation Research Part B: Methodological*, 46(9), pp.1239-1259.

Appendix A: The List of Top-28 Domestic Airports by Passenger Enplanements

In alphabetical order of the airports' three-letter codes: Atlanta, GA (ATL); Boston, MA (BOS); Baltimore, MD (BWI); Charlotte, NC (CLT); Washington, DC (DCA); Denver, CO (DEN); Dallas-Fort Worth, TX (DFW); Detroit, MI (DTW); Newark, NJ (EWR); Fort Lauderdale, FL (FLL); Dulles, VA (IAD); Houston, TX (IAH); Queens, NY (JFK); Las Vegas, NV (LAS); Los Angeles, CA (LAX); Queens, NY (LGA); Orlando, FL (MCO); Chicago, IL (MDW); Miami, FL (MIA); Minneapolis-Saint Paul, MN (MSP); Chicago, IL (ORD); Philadelphia, PA (PHL); Phoenix, AZ (PHX); San Diego, CA (SAN); Seattle-Tacoma, WA (SEA); San Francisco, CA (SFO); Salt Lake City, UT (SLC); Tampa, FL (TPA).

Appendix B: Shifts in Operations between NYC Airports

In 2017, both JFK and LaGuardia underwent runway reconstructions that temporarily reduced their air traffic capacity. If the FAA preemptively lifted slot control at Newark in order to allow the affected carriers to shift operations from JFK and LGA, our proxy variable for slot burning – usage of small aircraft in peak and offpeak slot periods – could be confounded by patterns of aircraft usage spilt over from JFK and LGA.

In order to test for evidence of spillover operations, we correlate the change in the frequency of scheduled flights by route between 2016 and 2017 and between 2016 and 2018.[25] We find no evidence of shifts in JFK's operations. The pairwise coefficients

[25] Figure 1 shows reduction in delays in 2018. This could be due to the fact that the reconstructed runways returned to operating at full capacity, or because it takes more than a year to shift operations between airports. For this reason, we study changes in scheduled flight frequencies in 2018 as well.

of correlation between the changes in scheduled flight frequencies at JFK and Newark are insignificant -0.0491 in 2017 and insignificant 0.0495 in 2018. However, we do find that the LGA routes that experienced a decrease in the number of scheduled flights in 2017 or 2018 (relative to 2016) tend to experience an increase in scheduled frequency at Newark, with the correlation coefficients of -0.2140 significant at 5% in 2017 and -0.2754 significant at 5% in 2018.

Airport	Newark		JFK		LaGuardia	
	Δ2017	Δ2018	Δ2017	Δ2018	Δ2017	Δ2018
Québec City, QC	216	472	-582	-1,156	–	–
Sarasota, FL	122	705	372	344	-685	-643
Jacksonville, FL	-84	220	6	70	-1,090	-1,316
Fort Myers, FL	593	556	-252	-531	-823	-1,093
Nantucket, MA	-34	156	-117	59	-20	262
Indianapolis, IN	395	889	559	1,029	-1,307	-1,054
Grand Rapids, MI	99	202	–	–	-169	348

Table 8: Change in the number of scheduled flights on routes that experienced a significant decrease at JFK and LGA and an increase at Newark.

We further investigate what routes experienced sizeable decrease in scheduled frequency (more than 15%) at JFK and LaGuardia and an increase in scheduled frequency at Newark. We identify seven such routes and document them in Table 8 above. None of these routes are a part of the sample of airports we use to test for slot burning, therefore we believe that possible shifts in operations did not affect patterns of aircraft usage at Newark in any spurious manner.

Effectiveness of Measures of Upward Pricing Pressure in Predicting Price Changes

Ratib M Ali

April 12, 2022

Abstract

I assess the effectiveness of the gross upward pricing pressure index (GUPPI) in predicting price changes of the 2013 merger between US Airways and American Airlines. I compute GUPPI using only publicly available data, and find that it is close to the observed average increase in price. However, unlike most markets, flights to/from Reagan Airport experience a price drop, likely due to mandated structural remedies; the GUPPI predicts a price increase at Reagan Airport, whereas a full merger simulation correctly predicts a price reduction.[1] I argue that the divergence between GUPPI and, if appropriate, the more accurate predictions of the merger simulation is due to the weaker assumptions made under the simulation. This underscores the fact that while GUPPI, with its restrictive assumptions and low computational burden, can be a good primary screening tool, it does not negate the necessity of employing a more rigorous secondary tool (such as a merger simulation) when assessing mergers.

I Introduction

Antitrust economists rely on measures of upward pricing pressure (UPP) to adjudicate whether a proposed merger will adversely affect consumers. Faster to

[1] The Ronald Reagan Washington National Airport (DCA) is located in Arlington, Virginia, near Washington, DC.

compute and with lower informational requirements than a merger simulation, UPP indices have gained traction in the antitrust community. However, the low informational and computational burden may result in reduced accuracy. This study uses retroactive merger analysis to investigate the ability of a measure of UPP to accurately predict post-merger prices in the 2013 US Airways-American Airlines merger.

Upward pricing pressure (UPP) "evaluat[es] potential unilateral effects in merger cases involving differentiated products" (Moresi, 2010). To illustrate, consider Firm A that wants to raise its prices. It fears that it will lose some of its consumers to Firm B. If, however, Firm A merges with Firm B, Firm A will be able to recapture the consumers who defect from Firm A to Firm B, because they are under joint ownership. Thus, the merger neutralizes the competitive threat from Firm B, increasing the payoff from raising Firm A's own price. This incentive results in an upward pricing pressure in the market. Measures of UPP are considered an improvement over concentration-based methods (such as using HHIs), and are thus frequently used as a screening tool for potential unilateral effects (Farrell and Shapiro, 2010). A popular measure of upward pricing pressure is the gross upward pricing pressure index (GUPPI), which measures the unilateral price effect due to the acquiring firm's ability to recapture sales diverted to the newly acquired firm.

This paper tests the reliability of the predictions made by the GUPPI in the context of the 2013 merger between US Airways (US) and American Airlines (AA). In this paper, I compute the predicted price changes implied by (i) the GUPPI and (ii) a merger simulation. I demonstrate circumstances under which the predictions differ, and explore the assumptions that lead to this difference. I find that the GUPPI sensibly identifies mergers that should be subject to further scrutiny, but fails to capture the reality of complex mergers in markets that include structural remedies. Specifically, the GUPPI reasonably approximates the observed average price change due to the 2013 merger between US Airways and American Airlines (7% vs. 8%). In contrast, a merger simulation better predicts outcomes in markets that involve structural remedies. Specifically, for routes involving Reagan Airport (the only airport to experience significant structural remedies), the GUPPI fails to predict the observed decrease in price (6% vs. −4%), whereas merger simulations provide results closer to reality (−0.3%). The structural remedies to the US Airways-American Airlines merger were divestiture of slots to low-cost carriers at Reagan National Airport, where a slot is a permission to perform one departure or one arrival by the airline at Reagan Airport. The merger simulation fares better than GUPPI because it makes weaker assumptions about rival responses. A summary of the results from each model with their underlying assumptions are provided in Table 1.

Table 1: Comparison of Models Applied to Reagan National Airport

Model	Change in price	Rivals respond to price	Divestment	Rivals choose network
Observed price change	−4.36%	Yes	Yes	Yes
GUPPI	5.66%	No	No	No
Merger simulation without divestment	3.06%	Yes	No	No
Merger simulation with divestment	−0.26%	Yes	Yes	No

There are several well-known limitations to the use of UPP indices. Calculations of GUPPI (or most other measures of upward pricing pressure) rely on pre-merger estimates of elasticities for post-merger price predictions. Mergers between two large firms, however, may violate this assumption due to improvements in product quality, or market segmentation by the merging firms. Furthermore, some measures of UPP assume a 10% efficiency gain, even though cost efficiencies may not materialize.[2] Thus, assuming a positive efficiency gain (or even assuming that the merger is efficiency-neutral) may err in favor of the merging firms.

UPP indices rely on diversion ratios, which measure the proportion of sales captured by a substitute product when the price of a focal product is increased. When diversion ratios are computed as aggregate measures, they may fail to capture the potential distributional impacts of a merger by missing the possibility of a region or submarket being underserved (e.g. due to a decrease in product choice, or a regional increase in price). To avoid missing these regional effects, diversion ratios and UPP indices may be calculated separately for geographically isolated markets. For example, this might comprise of a few square blocks for a grocery store (consisting of residents who mostly shop from one of a few grocery stores in the area), or half a state in the case of specialized hospital services. However, using geographically isolated markets undercuts the appeal of UPP indices, which were designed in part to avoid the messy affair of defining markets. Lastly, most UPP indices are meant to capture only unilateral effects between the merging parties. A merger between firms A and B may incentivize firm C to change its price if prices are strategic complements or substitutes, which can be captured by a correctly-specified structural model.

There are two main reasons to believe, *a priori*, that a full merger simulation

[2]Some reasons cost efficiencies may not arise include organizational chaos and loss of employee morale. The MIT Sloan Executive Education Blog notes that the success of airline mergers in general, and the US Airways-American Airlines merger in particular, depends on employee satisfaction (MIT, 2013).

from a structural model will better predict post-merger price effects than GUPPI. The first reason is its incorporation of rival responses. A merger simulation can, for example, predict Delta Airlines' price response due to the merger. GUPPI calculations assume no price response from Delta – a strong assumption, especially considering the entry of low-cost rivals due to the divestiture. The second reason is that the use of a full merger simulation permits one to model the impact of structural remedies on demand and prices.[3]

Modifying the GUPPI to account for the complexities of structural remedies all but eliminates the computational advantage it has over a full merger simulation. By highlighting the limitations of using GUPPI, I emphasize its strength – as a quick screening tool, and not a rigorous prediction of expected post-merger effects.

I.1 Literature Review

The limitations of UPP indices are well-documented in the literature. In this section, I highlight some of these limitations and discuss how a full merger simulation using a structural model might overcome them.

An important component in estimating measures of UPP is the diversion ratio. There are multiple methods to recover a diversion ratio. While Hausman (2010) argues that using a full structural demand system should be the only method of recovering diversion ratios, Farrell and Shapiro (2010) argue that firms' internal estimates of diversion ratios (tracked during their "normal course of business") can also be useful. The UK Competition & Markets Authority (2017) often relies on

[3]A third possible benefit of a structural model is its ability to endogenize medium-run investment decisions (for example, route choice in an airline merger); however, this is beyond the scope of the present paper because this calculation requires knowledge of fleet composition, which is proprietary.

consumer surveys soliciting second-choice data.

Conlon and Mortimer (2021) show that the diversion ratio is a structural parameter of each individual consumer, and that diversion measured after an increase in the price of a product recovers a local weighted average of the individual diversion ratios among all individuals who bought the product at the lower price but no longer buy the product at the higher price (an average treatment on the treated, ATT).

In contrast, a diversion ratio recovered using consumer surveys about second-choice data (à la CMA) is an average treatment on the untreated (ATUT), because all individuals who previously purchased the product (untreated) can no longer do so (treatment). Conlon and Mortimer (2021) show that different methods of estimating diversion ratios recover different treatment effects; this understanding is useful in interpreting diversion ratios estimated from different data sources.

Using the diversion ratio to approximate the proportion of consumers recaptured by a newly acquired merger partner is a simplification of many real-world complexities; for instance, it ignores the increased likelihood that remaining post-merger firms may change their own prices. A full merger simulation allows for the possibility that non-merging rivals may change their prices after the merger.[4]

UPP indices cannot predict a post-merger price without calibrating the merger pass-through rate, which may temper the merged firm's ability to increase prices following a merger. The merger pass-through rate is a measure comprised of local

[4]In the case of the US Airways-American Airlines merger, a full merger simulation also allows firms operating out of Reagan Airport to maximize profits subject to the slot constraint, accounting for slot allocation and divestiture.

second-order conditions on the demand function. While there are disagreements on whether the appropriate merger pass-through rate is estimated at a pre-merger or post-merger level, Jaffe and Weyl (2013) show that for a small GUPPI, pre- and post-merger pass-through rates are similar and can be approximated by first-order conditions. This allows antitrust economists to use pre-merger pass-through rates approximated using first-order conditions as the merger pass-through rate. A profit-maximizing price estimated from a utility model in which consumers respond to changes in product characteristics accounts for this pass-through effect without need for calibration.

UPP indices are used to screen for potentially anticompetitive mergers that warrant structural remedies. For instance, the UK competition authorities, as an operational rule, typically prescribe structural remedies for mergers crossing a UPP threshold (UK Competition & Markets Authority, 2017). However, a UPP index cannot measure the potential effects of proposed remedies. In order to measure the impact of the structural remedy on consumers, we need to employ a structural model where firms maximize profits subject to the prescribed remedies.

The rest of this paper is as follows: Section II explains the merger, along with the data available for the study; Section III outlines the most popular methods of computing upward pricing pressure, and explains the GUPPI in detail; Section IV compares and discusses the difference in predicted prices from a quick measure like upward pricing index with a merger simulation; Section V concludes.

II Background

Airline mergers have been common since deregulation in the 1970s. The use of slots as a structural remedy to mergers and airline partnerships has become commonplace in the last 15 years. Notable instances include the mergers between Delta Airlines and Northwest Airlines (2008), United Airlines and Continental Airlines (2010), US Airways and American Airlines (2013), Alaska Airlines and Virgin America (2016), as well as voluntary slot swaps between United Airlines and Delta Airlines (2015) and a joint venture between American Airlines and JetBlue Airways (2020).

Prior to the 2013 US Airways-American Airlines merger, the US airline industry was fairly concentrated, with five firms controlling about 80% of the market share. As is customary in airline mergers, antitrust authorities considered airport-pairs as relevant antitrust markets. The Antitrust Division at the United States Department of Justice (DOJ) along with some states successfully pursued divestment of airport slots and gates at Reagan Airport near Washington, DC, and LaGuardia Airport in New York City, NY, as pre-conditions for the merger in order to mitigate the negative effects of concentration at these airports.[5][6] The merger and the slot divestments were consummated in 2014.

I obtain data on prices and quantities from the publicly-available DB1B database (Bureau of Transportation Statistics, 2015), which I use to calculate market shares and diversion ratios. I use the publicly-available Schedule P.1.2 (Bureau of Transportation Statistics, 2012) to obtain data on airline markups. The purpose

[5] The co-plaintiff jurisdictions were the Attorneys General of Arizona, District of Columbia, Florida, Michigan, Pennsylvania, Tennessee, and Virginia.

[6] Other airports with small levels of slot divestment were Boston Logan Airport, Miami International Airport, Chicago O'Hare Airport, and Dallas Love Airport (United States v. US Airways Group & AMR Corporation, 2013). The State of Texas settled separately to ensure service to small rural communities.

of only using publicly-available data is to mimic the conditions that antitrust regulators would face when evaluating such mergers under deadlines imposed by the Hart-Scott-Rodino Act of 1976.

I limit my analysis to the 50 largest airports in the US by domestic passenger enplanements. I use only the third quarter to control for seasonality, using 2012 as the pre-merger period and 2015 as the post-merger period.[7]

III Theory

Of the prominent measures of upward pricing pressure, the Gross Upward Pricing Pressure Index (GUPPI) has the least informational requirement and makes the fewest assumptions (European Economic & Marketing Consultants, 2013). Other measures of upward pricing pressure include the confusingly named "upward pricing pressure" (European Economic & Marketing Consultants, 2013), which allows for positive efficiency gains from the merger. The efficiency gain is usually calibrated at 10% (or some other value based on industry consensus), and robustness checks are conducted. The "illustrative price rise" test, on the other hand, makes stronger assumptions about the demand function (van der Veer, 2012). As a result, I use GUPPI as the measure of upward pricing pressure throughout this paper.

III.1 Deriving the GUPPI

GUPPI measures the value of sales lost to the (former) rival as a proportion of the total revenue lost by the firm due to a price increase. The lost revenue to the former rival is recaptured by the merged entity following the merger, which creates an upward pricing

[7] Although the divestments were declared in late 2013, it took a number of months to consummate the transfers. This disqualifies 2014 as a candidate for post-merger data.

pressure. Mathematically, it can be expressed as (for example, see US Department of Justice and Federal Trade Commission (2010)):

$$GUPPI_{US,AA} = \frac{\text{Value of sales diverted to AA}}{\text{Revenues lost by US}} \tag{1}$$

The terms in the expression can be expressed in terms of quantities, prices, and costs, as follows:

$$GUPPI_{US,AA} = \frac{\text{Number of units diverted to AA} \times \text{Unit margin of AA}}{\text{Number of units lost by US} \times \text{Unit price of US}} \tag{2}$$

$$GUPPI_{US,AA} = \underbrace{\frac{\text{Number of units diverted to AA}}{\text{Number of units lost by US}}}_{DiversionRatio} \times \underbrace{\frac{\text{Unit margin of AA}}{\text{Unit price of AA}}}_{Markup} \times \underbrace{\frac{\text{Unit price of AA}}{\text{Unit price of US}}}_{PriceRatio} \tag{3}$$

The first term in equation 3 is called the diversion ratio from US Airways to American Airlines (discussed in Section III.2). The second term refers to the percentage markup of American Airlines. Thus, GUPPI can be expressed as follows:

$$GUPPI_{US,AA} = DiversionRatio_{US,AA} \times Markup_{AA} \times PriceRatio_{AA,US} \tag{4}$$

In words, the upward pricing pressure is a function of the quantity of passengers who would divert from US Airways to American, American's price markup, and the relative prices of the two products. The greater the proportion of consumers that would defect from US Airways to American Airlines, and the greater the price markup, the greater is the upward pricing pressure.

III.2 Diversion Ratio

A crucial component of any measure of UPP is the diversion ratio, "which measures the fraction of consumers that switch from one product to an alternative after a price increase" (Conlon and Mortimer, 2021). Ordinarily, diversion ratios can be inferred from sales data or internal business records (Farrell and Shapiro, 2010), although Hausman (2010) insists on a full structural model.

The simplest measure of diversion ratio from US Airways to American Airlines is the number of passengers diverted from US Airways to American, as a proportion of passengers who choose a non-US Airways flight due to a change in the price of flights by US Airways. Mathematically, for a change in the price of US Airways,

$$DiversionRatio_{US,AA} = \frac{\partial Q_{AA}}{\partial P_{US}} / \mid \frac{\partial Q_{US}}{\partial P_{US}} \mid \tag{5}$$

Own-price and cross-price elasticities of demand ($\varepsilon_{US,own}$ and $\varepsilon_{US,AA}$) for a change in the price of US Airways are defined as:

$$\varepsilon_{US,own} = \frac{\partial Q_{US}}{\partial P_{US}} \times \frac{P_{US}}{Q_{US}} \tag{6}$$

$$\varepsilon_{US,AA} = \frac{\partial Q_{AA}}{\partial P_{US}} \times \frac{P_{US}}{Q_{AA}} \tag{7}$$

Therefore, the diversion ratio can be expressed in terms of the elasticities. For a change in the price of US Airways:

$$DiversionRatio_{US,AA} = \frac{Q_{AA}}{Q_{US}} \times \mid \frac{\varepsilon_{US,AA}}{\varepsilon_{US,own}} \mid \tag{8}$$

Under the assumptions of a multinomial logit discrete choice utility model (including the IIA assumption), the elasticities can be expressed as:

$$\varepsilon_{US,own} = \alpha \times (1 - share_{US}) \times price_{US} \tag{9}$$

$$\varepsilon_{US,AA} = \alpha \times price_{AA} \times share_{AA} \tag{10}$$

where α is the price coefficient from the logit model. Thus, the diversion ratio simplifies to the following:

$$DiversionRatio_{US,AA} = \frac{Q_{AA}}{Q_{US}} \times \frac{\varepsilon_{US,AA}}{\varepsilon_{US,own}} \tag{11}$$

$$DiversionRatio_{US,AA} = \frac{Q_{AA}}{Q_{US}} \times \frac{\cancel{\alpha} \cdot price_{AA} \cdot share_{AA}}{\cancel{\alpha} \cdot (1 - share_{US}) \cdot price_{US}} \tag{12}$$

In Equation 12 above, I derive an expression for diversion ratio for the commonly used logit discrete choice model of demand (Anderson et al., 1992). The diversion ratio can be computed using only publicly available data of prices and quantities, and mathematically analogous to deriving the ratio using elasticities from a structural model (that assumes logit discrete choice demand).

IV Results

IV.1 Observed Price Increase After Merger

Table 2 outlines the price change due to the merger for markets including the 50 largest airports, and only markets involving Reagan Airport. The table shows that while prices on average went up by 7.6% due to the merger, the price fell by 4.4% for routes involving Reagan Airport. Special attention is given to Reagan Airport due to the significant structural remedies imposed by the DOJ – DOJ required divestment of about 20% of all slots held by the newly merged entity (15% of all slots) to low-cost carriers (Ali, 2020).

Table 2: Observed increase in price following US-American merger

Year	Pre-merger price 2012	Post-merger price 2015	Pct increase
Price for mkts involving:			
Top 50 airports	252.58	271.77	7.60%
	(59.13)	(57.46)	
Reagan Airport	263.86	252.36	-4.36%
	(67.71)	(50.96)	

Parentheses contain standard deviation.

IV.2 GUPPI calculation

Equation 4 above outlines how to calculate GUPPI. Diversion ratio is computed using Equation 12 referenced above. Price and market shares for the two airlines are computed for each market using the publicly available DB1B database. The quantities used in the equations above are the number of passengers flown by the respective carriers within the market, rendering a single diversion ratio for each market. A markup of 12.08% is imputed from the publicly available Schedule P.1.2 database. Section IV.3 discusses the markup calibration in greater detail and performs robustness checks. Following Berry and Jia (2010), I define a market as a non-directional city-pair in a given quarter. Table 3 shows that for the top 50 airports in 2012, the GUPPI for the US-American merger is 6.6%.

From Table 3, I find that the GUPPI (6.6%) closely resembles the observed increase in price (7.6%). In that, the GUPPI is adequate in predicting the general increase in price due to the merger. Unlike GUPPI, the observed increase in price includes secular macroeconomic trends. In Section IV.4, I compare the GUPPI

Table 3: Calculation of GUPPI for routes involving 50 busiest airports

	US	American
Price	252.6908	252.1912
	(55.63124)	(75.13782)
Quantity	742.3446	366.3041
	(635.9868)	(675.46)
Market share	0.3357	0.0977
	(0.2409)	(0.1562)
Markup		0.1208
Diversion ratio	0.4810	
	(1.8133)	
GUPPI	6.58%	
	(0.2994)	
Observed price increase	7.60%	

Parentheses contain standard deviation.

computed only for markets involving Reagan Airport with the simulated merger, neither of which take the time trend into account. As such, that the observed average increase in price includes macroeconomic trends is not of consequence, since the comparison is between methods that both include macroeconomic trends.

IV.3 Calibration of Markup and Robustness Check

Accounting based measures of markups often differ from true economic markups because accounting based measures of costs include sunk costs, exclude opportunity costs, and fail to distinguish average costs from marginal costs. The standard approach for estimating economic markups relies on first estimating a structural demand model in order to infer marginal costs from firm first-order conditions. To keep my calculated GUPPI computationally simple, however, I measure markups using publicly available accounting measures available with the Bureau of Transportation Statistics. Column [1] in Table IV.3 lists the reported revenue earned from transporting passengers. Column [2] in Table IV.3 lists the carrier's pre-tax income from all sources, which includes accounting costs. The calibrated markup of 12.08% for American Airlines is within the same order of magnitude as other industry documents.[8] Markups for various carriers for the third quarter of 2012 are provided below from Bureau of Transportation Statistics (2012):

Table 4: Markup for airlines, third quarter of 2012

Airline	Passenger Revenue ($ thousands) [1]	Pre-tax Income ($ thousands) [2]	Markup [2]÷[1]
American Airlines Inc.	1,087,686	131,414	12.08%
Delta Air Lines Inc.	430,524	85,435	19.84%
JetBlue Airways	227,137	61,860	27.23%
Southwest Airlines Co.	38,023	329	0.87%
Spirit Air Lines	31,412	7,287	23.20%
United Air Lines Inc.	627,416	46,787	7.46%
US Airways Inc.	165,014	19,650	11.91%

In Table 5, I perform robustness checks by varying the calibrated value of the markup. I find that the results are sensitive to our calibration of the markup. The robustness check was conducted using the industry average for markups (9%), and an arbitrarily chosen high markup, which is still possible for some airlines during highly profitable periods. GUPPI fluctuates between 4.9% for a low markup (9%) and 10.9% for a high markup of 20%. This should not be taken to dismiss our value of GUPPI, since the calibration was done using credible information, but rather as a reminder that proper calibration is required to obtain accurate results.

Table 5: Calculation of GUPPI with varying markups

	(1)	(2)	(3)
Markup (calibrated)	12.08%	9%	20%
GUPPI	6.58%	4.90%	10.90%
Std. Dev.	0.2994	0.2231	0.4957

[8]For instance, a *The Wall Street Journal* article puts the average 2018 airline markup at 9%.

IV.4 GUPPI for Reagan Airport

Next, in Table 6, I run the same calculations for the diversion ratio and GUPPI, but only for routes involving Reagan Airport. DOJ required divestment of about 20% of all slots held by the newly merged entity (15% of all slots) to low-cost carriers, which is significant but should not be reflected in our measure of upward pricing pressure.

Table 6: Calculation of GUPPI for routes involving Reagan Airport

	US	American
Price	266.4815	249.3104
	(61.4244)	(67.1112)
Quantity	1131.136	284.4984
	(795.4998)	(537.1485)
Market share	0.5091	0.0909
	(0.2525)	(0.1683)
Markup		0.1208
Diversion ratio	0.4123	
	(1.4908)	
GUPPI	5.66%	
	(0.1986)	
Observed price increase	−4.36%	

Parentheses contain standard deviation.

Table 6 shows a GUPPI of 5.7%, which is different from the observed price *decrease* of 4.4%. Indeed, nothing in the set-up of the formula to calculate GUPPI [equation (4)] can incorporate changes in the market structure, or even consider a decrease in price following a merger. While the GUPPI (5.7%) does not adequately predict the price change for routes involving Reagan Airport (−4.4%), the price change offered by a merger simulation (−0.26%, discussed in Section IV.5) is closer to the observed price decrease, and is negative.

IV.5 Merger Simulation

I begin by using the structural model developed in Ali (2020) on the pre-merger data for Reagan Airport. The restriction on Reagan Airport allows me to focus on markets that experienced change in market structure. By restricting the dataset to the pre-merger time period (third quarter of 2012), I am only using information available to the econometrician at the beginning of 2014, and using the same data that will be used to compute the GUPPI. These restrictions allow me to directly compare the results from the merger simulation with the GUPPI.

The estimates from the structural model in Table 7 are consistent with our prior understanding of air travel. All travelers derive disutility from price and utility from more frequent flights. Leisure travelers are more sensitive to changes in price (coefficient on price = -0.0262, elasticity = 6.0573) than business travelers (coef = -0.0049, elasticity = 1.1259), whereas business travelers are more sensitive to changes in frequency (coefficient on frequency = 0.0420) than leisure travelers (coef = 0.0295). The model predicts that 38% of all passengers travel for business.

I then use this model to run two simulations. In both simulations, I need to make an assumption about the endogeneity of airline networks. Airline networks are generally assumed to be pre-determined (see, for instance, Berry and Jia (2010)) due to long-negotiated labor, hotel, refueling, and hangar contracts. This assumption is strong when it comes to a merger, because the time needed to consummate a merger is long enough to renegotiate such contracts, and is a potential source of efficiency gains. As a compromise between assuming a fixed network and a fully re-optimized network after the merger, I assume that the newly merged entity eliminates duplicate/rival products previously offered by the acquired airline on

Table 7: Demand Coefficients

Covariate	Type-Business	Type-Leisure
	Business	Leisure
Constant	-3.2838	0.1016
	(1.4353)	(0.3655)
Layover	-18.5045	-1.6016
	(0.1089)	(0.1280)
Tourist	-0.2294	-1.5230
	(0.0002)	(0.2195)
Logdist	-0.4344	-1.3451
	(0.0191)	(0.0023)
Closest	0.0006	0.0002
	(0.0067)	(0.0584)
Legacy	-3.2739	-4.7257
	(0.0179)	(0.0243)
SD price	0.0221	0.0187
	(0.0142)	(0.0000)
Price	-0.0049	-0.0262
	(0.0083)	(0.0000)
Frequency	0.0420	0.0295
	(0.0001)	(0.0001)
Proportion*	0.3811	0.6189
	(0.0002)	

Parentheses contain standard errors.

* The proportion of leisure-type travelers is the complement of proportion of business-type travelers, and thus do not have standard errors of its own.

routes previously served by both airlines.

In the first simulation, I simply merge the ownership of the slots owned by American and US Airways to one entity, without any divestments. Each product (flight) owned by this new entity inherits the product characteristics of the original product, including its product unobservables, ξ, estimated using the structural model. In order to simulate the effects of a merger, in markets where both American and US Airways competed, I drop all products offered by American Airlines.[9] As

[9]Although the merged entity kept the "American" brand, the merger was essentially US Airways acquiring American Airlines. The management of the new entity came from US Airways, and US Airways was the airline that maintained a hub (and the dominant market share) at Reagan Airport.

shown in Table 8, my simulation predicts a 3.06% increase in price following the merger with no divestment for markets involving Reagan Airport.

The second simulation simulates a merger with the prescribed divestment. I again assume that the products being offered by the new entrants retain the product characteristics of the pre-merger product. My simulation, shown in Table 8, predicts a 0.26% fall in prices following the merger with divestment.

Table 8: Price increase at Reagan Airport, as predicted by merger simulation

	(Qty-weighted) mean price	Difference from actual price	Percentage difference from actual price	
Pre-merger price	263.86			
Observed post-merger price	252.36	−11.51	−4.36	%
Merger without divestment	271.93	8.07	3.06	%
Merger with divestment	263.18	−0.68	−0.26	%

A caveat of this study is that the simulations assume the network as exogenous. In other words, American is not allowed to change its network in response to entry by JetBlue. This restrictive assumption is made to avoid the computational burden imposed by endogenizing entry. However, this simulation allows for firms to competitively adjust their prices, which is an improvement over the GUPPI.

I find that a merger simulation (−0.26%, shown in Table 8) performs better than calculation for GUPPI (5.7%, shown in Table 6) in predicting the price change due to merger (actual price change was −4.4%, shown in Table 8). There are two main reasons for this divergence in the two methods. Firstly, GUPPI does not account for changes in efficiency or shifts in demand after the merger. Other screening tools used instead (or in conjunction) of GUPPI, like the UPP, assume an *increase* in efficiency. Thus, GUPPI assumes mergers to be efficiency-neutral, whereas UPP assumes positive gains in efficiency. A merger simulation, on the other hand, allows for shifts in demand by making weaker assumptions on product quality after the merger, which can result in a decrease in price. In a merger simulation, the post-merger product unobservable characteristic, ξ_j, can be a proxy for changes in efficiency and product quality (manifested through a change in mean product demand). In this paper, my choice of the post-merger product unobservable assumes the merger is efficiency-neutral, the same assumption made by GUPPI. In Ali (2020), my choice of the range of product unobservables from the data imply the possibility of an efficiency gain or loss.

Secondly, the assumptions underlying GUPPI do not account for competitive responses by rivals. Prices can be strategic complements or substitutes; a merger simulation allows firms to optimally set price, which is not the case for measures of upward pricing pressure like the GUPPI.

Neither the GUPPI nor the merger simulation detrends for macroeconomic trends. While the merger simulation provides a price change that could be considered as inaccurate, it performs better than GUPPI, and can produce negative values. This counterexample highlights GUPPI's inability to predict negative price changes, which is not a limitation of the merger simulation. Therefore, the GUPPI may not always be adequate for predicting price changes due to a merger, especially ones with structural remedies or other changes in market structure.

V Conclusion

This paper finds that for the simplest cases (that is, one without significant changes in market structure), measures of upward pricing pressure may approximate the order of magnitude of the change in price. The quality of the approximation will largely depend on the accuracy of the input variables, which may be easier to find in some industries than others. However, such measures should only be used as a screening tool, and not a conclusive one.

For mergers where structural remedies are contemplated, a merger simulation is necessary to understand the impact of the merger with remedy. A merger simulation can incorporate changes in the market structure due to the proposed remedy and strategic decisions made by firms not involved in the merger. I show that a merger simulation can predict decreases in price, which GUPPI (and other indices of upward pricing pressure), by design, cannot. This is due to the restrictive assumptions made by GUPPI – that mergers are efficiency-neutral, that rivals do not competitively respond to the merger, and that the market structure remains indentical after the merger.

Without undertaking such analyses, the act of proposing structural remedies becomes an arbitrary exercise. For instance, in the absence of merger simulation, the question of whether 104 slots at Reagan Airport should have been divested as opposed to any other number, and whether these slots should have exclusively gone to low-cost carriers, would not be based on econometric estimations or quantitative analysis informed by an underlying economic model of firm behavior.

The results underscore the need for conducting rigorous pre-merger analyses before approving mergers or while designing remedies.

References

Ali, R. M. (2020). Effect of Airport Slots in Competition and Antitrust Policy: Evidence from a Recent Merger. Working Paper.

Anderson, S. P., de Palma, A., and Thisse, J.-F. (1992). Discrete Choice Theory of Product Differentiation. The MIT Press.

Berry, S. and Jia, P. (2010). Tracing the Woes: An Empirical Analysis of the Airline Industry. American Economic Journal: Microeconomics, 2(3):1–43.

Bureau of Transportation Statistics (2012). Schedule P.1.2.

Bureau of Transportation Statistics (2015). Airline Origin and Destination Survey (DB1B).

Conlon, C. T. and Mortimer, J. H. (2021). Empirical Properties of Diversion Ratios. Rand Journal of Economics, 52(4):693–726.

European Economic & Marketing Consultants (2013). UPP, GUPPI and IPR – Merger Screening Tools .

Farrell, J. and Shapiro, C. (2010). Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition. The B.E. Journal of Theoretical Economics, 10.

Hausman, J. (2010). 2010 Merger Guidelines: Empirical Analysis. Working Paper.

Jaffe, S. and Weyl, G. (2013). The First-Order Approach to Merger Analysis. American Economic Journal: Microeconomics, 5(4).

MIT (2013). The role of employee satisfaction in successful airline mergers. MIT Sloan Executive Education Blog.

Moresi, S. (2010). The Use of Upward Price Pressure Indices in Merger Analysis. The Antitrust Source.

UK Competition & Markets Authority (2017). Retail mergers commentary.

United States v. US Airways Group & AMR Corporation (2013). United States District Court for the District of Columbia, 13-cv-1236.

US Department of Justice and Federal Trade Commission (2010). Horizontal merger guidelines.

van der Veer, J. P. (2012). UPP – frequently asked questions.