Structural Remedies in Network Industries: An Assessment of Slot Divestitures in the American Airlines/US Airways Merger*

Yongjoon Park†

[Job Market Paper]

November 25, 2019

The latest draft is available here

Asset divestitures are often negotiated to alleviate anticompetitive concerns created by horizontal mergers. I develop a structural model to evaluate the effectiveness of alternative slot divestiture schemes in the US airline industry, focusing on the divestiture of slots at Ronald Reagan Washington National Airport (DCA), which the government required as a condition of the American/US Airways merger. Departing from the existing literature, my model accounts for how the number of slots allocated to a route segment affects carrier costs, how passengers going to many different destinations may use the same segments, and how carriers choose to allocate slots to segments. I use counterfactuals to show that slot divestitures can result in the re-allocation of surplus between consumers; to estimate the proportion of slots that the merged American would have needed to divest to maximize total welfare; and, to evaluate the effects of allocating divested slots to different types of carriers. I find that the proposed divestiture raised consumer surplus significantly ($112M per year) compared to approving the merger without divestiture, but that it re-allocated surplus between consumers in different markets. I also find that the policy of only allowing the slots to be divested to low-cost carriers raised consumer surplus relative to the policy of only allowing the slots to be divested to legacy carriers.

Keywords: merger remedy, slot divestitures, endogenous entry, horizontal merger analysis, static games, airlines, network effect

JEL Codes: C35, C54, L44, L13, L93

*I am very grateful to my advisor Andrew Sweeting for his invaluable guidance and support. I also thank Ginger Jin, Dan Vincent, John Shea, Joshua Linn, Guido Kuersteiner, Christy Zhou, Chenyu Yang, Yong Paek, Xinlu Yao for their helpful comments. This paper is previously circulated as "The Welfare Implications of Slot Divestitures." All errors are my own.

†Department of Economics, University of Maryland. Email: joon@umd.edu

1 Introduction

When large horizontal mergers are proposed, a number of standard approaches such as merger simulations and upward pricing pressure can be used to predict whether or not the merger may reduce competition. In cases when the merger may make select markets less competitive, antitrust authorities negotiate a set of remedies with the merging parties, rather than blocking the merger outright.[1] The preferred form of merger remedy is a structural remedy (asset divestiture), designed to redistribute assets to market rivals or new entrants in a way that they can restore competition that could be lost through the merger—see United States (2019).[2] These assets may include production capacity, distribution centers, wireless spectra, airport slots or intellectual property. However, it is hard to apply the standard approaches to evaluate merger remedies because the link between these types of assets and their costs or the set of assets offered to consumers is not specified.

In this paper, I develop a model of the domestic airline industry that can be used to evaluate alternative take-off/landing slot divestitures that may be proposed for airline mergers. I use the model to examine the slot divestitures imposed on the merger between American Airlines and US Airways (AA/US) at Ronald Reagan Washington National Airport (DCA); however – as I emphasize in the conclusion – this model could also be leveraged to evaluate a number of other slot-related policies.[3] As is the case at New York airports, carriers at DCA are required to hold a landing slot for each scheduled departure or arrival, with a cap on the number of slots each hour in order to mitigate congestion. AA/US transported a combined 60% of DCA passengers and a 56% share of the departures and arrivals pre-merger. Given the dominant position of both firms and the slot constraints at DCA, the U.S. Department of Justice (DOJ) argued that the slot constraints would prevent rivals from expanding or initiating their service if the merged American (NewAA) raised prices—see United States (2013). To gain approval for the merger, AA/US agreed to divest 104 slots (approximately 15% of its holdings) to low-cost carriers (LCCs), as well as forfeit their rights and interests at any associated gates or

[1] Merger remedies have become a dominant method for resolving merger-related competition concerns. For example, of the merger cases challenged by U.S. antitrust authorities between 2008 and 2016, eight times more cases were resolved with settlement/remedies than with court proceedings—see Hatzitaskos et al. (2019).

[2] Behavioral remedies (i.e. conditions imposed to mitigate or prohibit specified anticompetitive conduct) represent an alternative merger remedy. Examples of behavioral remedies include non-discrimination provisions, anti-retaliation provisions, and prohibitions on certain contracting practices. However, as behavioral remedies are often imposed against the merged firm's profit-maximizing incentives, these remedies are often less preferred than structural remedies. Additionally, most jurisdictions, including those in the U.S., have stated a preference for using asset divestitures over behavioral remedies in merger cases that present anti-competitive concerns —see Kwoka (2017).

[3] Examples are the competition effects of exchanges or sales of landing slots within an airport (Reitzes et al. (2014)), slot swaps, or alternative mechanisms for controlling congestion (Ball et al. (2007)).

other ground facilities at the airport. I use my model to quantify how divestitures may affect the welfare of passengers in different markets (defined as origin-destination pairs) and predict how passenger welfare may be affected by different divestiture decisions—for instance, more or less slots, or divestitures where slots are allocated to legacy carriers instead of LCCs (a policy that legacy rivals, such as Delta, support).

The model that I develop extends both the existing frameworks for merger simulation and the existing models of airline markets in several important ways. A key contribution is that I model the connection between divested assets (landing slots) and both carrier costs and how the allocation of assets to products affects pricing and costs. This type of connection has been missing in the existing merger simulation models. In the context of airline merger cases, divested assets are landing slots that a carrier may choose to use on a wide range of possible route segments. Because many passengers make connections, a single slot may be used to serve passengers in many different origin-destination markets.[4] My model takes into account this industry feature and enables cross-market interactions. This model feature is an extension to the existing entry models in airline markets (Berry (1992), Ciliberto and Tamer (2009), Ciliberto et al. (2018), Li et al. (2019)) where individual markets are typically assumed to be independent and there are no interactions across markets. Last, the existing airline entry models ignored the differences between flight segments (nonstop origin-destination pairs where carriers allocate capacity) and passenger markets, and they abstracted away from carriers' capacity choices. My model, on the other hand, allows carriers to endogenously choose their capacity, by distinguishing flight segments and markets.

I model a two-stage game. In the first stage, carriers choose which flight segments to serve, while in the second stage, carriers choose how many slots to allocate to each segment and set product prices in each market where their chosen selection of segments allows them to serve. I allow for increases in marginal costs associated with a carrier's load factor (passenger to seat ratio) on a particular segment. This implies that the allocation of capacity to a segment affects costs in a large number of passenger markets (for example, when American allocates more capacity to Raleigh-DCA, this will affect its costs of serving passengers traveling from Raleigh to DCA, from Raleigh to Hartford via DCA, Raleigh to Boston via DCA, etc). Key parameters of the model are estimated by using quarterly data taken from the publicly-available T-100 and DB1B databases from 2012-2013. In a range of counterfactuals through which I re-allocate slots across carriers, I use the estimated model to resolve for route selection, slot allocations, prices and welfare.

[4] Slot or gate divestitures have been the remedy that policy makers have imposed in several completed mergers (e.g. IAG/Aer Lingus in 2015, American/US Airways in 2013, United/Continental in 2011, Ryanair/Aer Lingus in 2006, and Air France/KLM in 2004). Li et al. (2019) use a more standard merger simulation approach to consider the effect of a remedy proposed in the failed 2001 merger between United and US Airways in which American offered to guarantee that it would serve specific routes.

The counterfactual analysis shows that requiring American Airlines and US Airways to divest slots raised aggregate consumer surplus, but harmed a subset of DCA passengers. Compared to the case where the merger was allowed with no divestitures, the divestiture of 15% of American's slots is estimated to have raised total consumer surplus by $28M per quarter (roughly, $112M per year). I estimate that slot divestitures would have raised consumer surplus, although total surplus increased little. However, these divestitures caused American to eliminate service on some marginally profitable routes, and the model predicts that the carriers that are allocated the slots would be unlikely to serve either these eliminated routes or the routes where American and US Airways were both active before the merger (overlapped routes) and for which competitive concerns are likely to be the greatest (e.g. DCA-Raleigh and DCA-Nashville markets). I therefore consider alternative divestiture policies through which the recipients of the divested slots are either required to use them to serve particular routes or the slots are granted to different carriers whose incentives to serve different types of routes might differ from the LCCs that were actually granted the slots. I find that requiring the slot recipients to serve the overlapped routes can create a net surplus gain. Additionally, I find that legacy carriers are likely to serve routes based in large cities (e.g. Miami, Boston) and are less likely to serve those routes based in small communities to which they expressed their intention to expand services.

Before discussing the related literature, I highlight some of the simplifications I make to maintain tractability or to make use of publicly-available data. First, DCA slots are tied to specific hours of the day, and there are some differences across slots depending on whether the slots can only be used for small planes and whether they can be used on routes outside the usual 1,250 mile range allowed from DCA (perimeter rule). Slots in my model, however, are assumed to be not time specific, single type, and used only for routes within the perimeter rule. One airport slot in the model is interpreted as the average number of operations per day on a particular directional segment. Second, the divestitures involve gates and other ground facilities, as well as slots. My analysis treats slots as the only relevant asset. Third, while I consider discrete choices to serve particular routes, I treat slot allocation as a continuous choice. Fourth, while I consider a relatively rich set of connecting products, I perform some aggregations so that the number of prices that I have to solve for is not too large and overly burdensome. Fifth, I also ignore the important heterogeneity in the willingness of different types of customers (e.g. leisure, business and government passengers) to substitute across products in response to changes in the availability of connections or prices. Finally, the counterfactuals may suffer from a multiplicity of equilibria. I discuss how I treat this issue below.

Related Literature

This paper contributes to the literature on merger and merger remedies, as well as the empirical literature on market entry and airline competition. In the literature, there are a large number of merger retrospective studies but those mainly focus on general individual merger effects rather than merger remedies. Despite the recent growing interests in tightening merger policy, attempts to systematically analyze merger remedies *ex ante* are rare in the literature.

The U.S. airline industry has experienced a large number of mergers since deregulation in 1978. Ashenfelter et al. (2014) summarizes studies on the impact of airline mergers on price. Most studies find price increases after mergers, but the magnitudes of those effects depend on the sample selection and empirical strategies employed—see Borenstein (1990); Kim and Singal (1993); Kwoka and Shumilkina (2010). The studies of recent airline mergers provide mixed results. While Hüschelrath and Müller (2014) suggests that the merger between Delta Airlines and Northwest Airlines increased short-run prices by 10%, Israel et al. (2013) claims that due to the network expansion of a merged firm, post-merger price increases may come from increases in a consumer's willingness to pay. However, the existing literature has not focused specifically on changes in market competition at slot-controlled airports even though it is at these airports where the Department of Justice has implemented remedies.

Studies that evaluate merger remedies are rare in the literature, and those that exist are skewed to *ex post* analysis. The Federal Trade Commission (FTC) conducted two studies, one in 1999 and another in 2017, that examine the efficacy of the merger remedies that the authority had ordered over the past two decades by closely analyzing the survival of divested assets—FTC (1999) and FTC (2017). They find that 70% of such divested assets remained in relevant markets and restored market competition; however, there has been criticism that the U.S. antitrust authorities are far too willing to accept remedies and fail to restore competition—Kwoka (2014). Similar analyses have been conducted in the context of other jurisdictions, including the EU, the UK, and Canada—see Duso et al. (2007); Wang and Rudanko (2012). While most merger remedy studies in this literature focus on descriptive analysis of what happen after remedies were imposed, my paper provides an *ex ante* assessment to measure the effectiveness of a merger remedy with respect to market competition and passenger welfare.

Empirical entry models in the literature can be largely categorized into two groups. Without modeling equilibrium price competition, most of the early literature (Bresnahan and Reiss (1990), Seim (2006) and, applied to airlines, Berry (1992), and Ciliberto and Tamer (2009)) allowed firms to make discrete entry decisions in independent markets. In contrast, Reiss and Spiller (1989), Eizenberg (2014), Fan and Yang (2016), Wollmann

(2018), Ciliberto et al. (2018), and Li et al. (2019) use two-stage models in which firms enter and then make explicit price choices. My model extends this second school of thought, allowing marginal costs to vary with load factors, interactions between multiple product markets, and choices of capacity, as well as price choices. I follow Eizenberg (2014), Fan and Yang (2016) and Wollmann (2018) in assuming that firms only learn product-specific demand and marginal cost unobservables after they have chosen which routes to serve. This choice allows me to estimate the fixed costs associated with serving a market separately from the demand and marginal cost functions. In contrast, Ciliberto et al. (2018) and Li et al. (2019) allow for demand and marginal cost shocks to be known when entry choices are made, which increases the computational burden. Extending the model to allow for selection that arises under complete information is one obvious direction for future research.

There is an airline paper that does allow for elements of network competition—see Aguirregabiria and Ho (2012). Since the paper assumes that each carrier has a local manager on each route who maximizes his or her local profit, the computational burden associated with network choices is reduced. In contrast to the study, I allow carriers to optimize entry, capacity and price choices over all of the segments and routes that they might serve, while focusing on a single airport of particular interest to reduce the computational burden.

Section 2 provides institutional background on Ronald Reagan Washington National Airport (DCA), a slot-controlled airport, as well as on slot divestitures. In Section 3, I develop a structural entry model. Section 4 describes the data used in this study. Section 5 discusses the estimation of the parameters. Section 6 reports the estimation results and model fit, and Section 7 provides the counterfactual analysis. Finally, Section 8 concludes the paper.

2 Institutional Background

This section discusses the concept of airport landing slots, slot-controlled airports, and slot divestitures in the context of the AA/US merger.

2.1 Slot-Controlled Airports

In aviation, an airport landing slot refers to operational authorization to either take off or land at a particular airport on a particular day during a specific time period. For example, if an airline carrier holds a Monday 7:00-7:59AM landing slot at an airport, the carrier has the right to take off or land during the time slot at that airport. While runway access at most U.S. airports is on a first-come-first-served basis and those airports do not use any slot controlling system, a few so-called *slot-controlled airports* in the U.S. regulate the maximum number of flights per hour to mitigate heavy congestion and delays, implying that slots at those airports are scarce and are valuable assets to carriers.

In the context of U.S. aviation, the first slot control system, the High Density Rule (HDR), was instituted by the Federal Aviation Administration (FAA) in 1968, and the cap on the number of hourly arrivals and departures was applied to five "high-density traffic airports"—John F. Kennedy (JFK), LaGuardia (LGA), and Newark (EWR) serving New York City; Ronald Reagan Washington National Airport (DCA) in Washington D.C.; and O'Hare (ORD) in Chicago. The FAA initially assigned slots to the carriers that already had them under scheduling committee agreements used prior to the HDR (i.e., "grandfather rights"). As slots are expected to be actively used, there is a "use-or-lose" provision that states that any slot not utilized 80 percent of the time over a period shall be recalled by the FAA. Generally, as long as airlines comply with the rules, they may continue to keep those slots.

While carriers can lease or trade their slots under FAA approval, secondary market slot sales are rare and prohibited at some airports. Trading slots is fairly common to facilitate airline schedules, and leasing slots is an attractive option for slot holders since they can control slots. However, the leasing of slots tends to be based on short-term agreements with early termination clauses, and leasing to new entrants is rare, as the entrants are considered direct competitors. A secondary market to buy/sell slots was created under the HDR in late 1980, but slot transactions were infrequent and airports other than DCA have not been authorized to buy or sell slots since early 2000, as the HDR was phased out and new slot control system was introduced at those airports.[5]

There are other airport-specific slot restrictions. First, a subset of airport slots at slot-controlled airports is designated only for small airplanes. At DCA, for example, the maximum number of flights per hour is 48, and 11 of them are designated only for commuter aircraft—aircraft that may be used only for operations with turboprop and reciprocating engine aircraft with no more than 76 seats or turbojet aircraft with fewer than 56 seats. This restriction was introduced to balance maximizing the economic use of runway resources and preserving services to smaller communities. Second, nonstop flights from/to DCA (LGA) are not allowed to exceed 1,250 (1,500) nautical miles, which we call the "perimeter rule." For the purpose of reducing airport congestion and inducing passengers to use alternative airports, the FAA has imposed perimeter restrictions on DCA and

[5] More detailed information is available at the Federal Aviation Administration (FAA) website: https://www.faa.gov/about/office_org/headquarters_offices/ato/service_units/systemops/slot_administration/

Figure 1: DCA Nonstop Flights and the Perimeter Rule

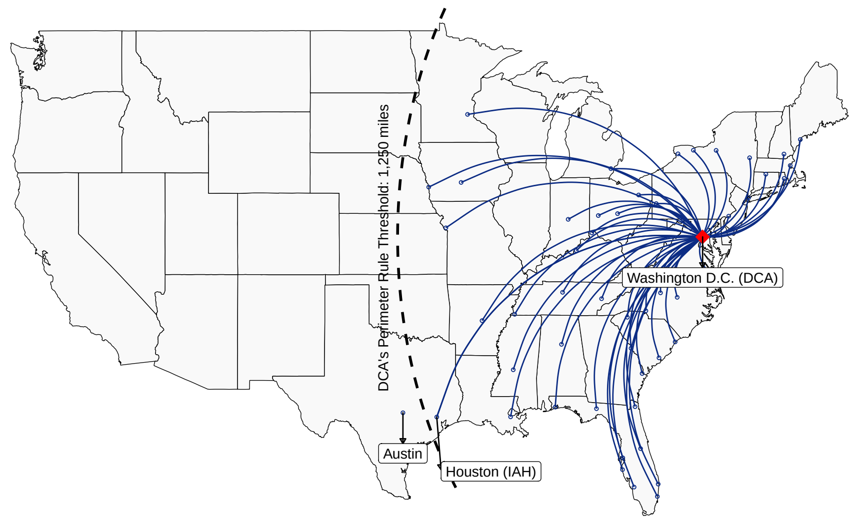

Note: The 1,250 nautical mile perimeter from DCA is shown as a curved dotted line. As an illustration of the perimeter rule, the distance between DCA to Houston airport is within the boundary, while the distance between DCA to Austin airport is not.

LGA. Figure 1 visualizes the set of nonstop flights to/from DCA within the perimeter rule. There are a few slots at the airport specifically designated for routes beyond the perimeter rule, such as DCA to Los Angeles (operated by American) and DCA to Denver (operated by United). Those designated slots are exempted from the airport's slot regulation, and carriers are required to obtain special authorization to operate services on those routes.

2.2 Slot Divestitures of the AA/US Merger at Reagan National Airport

Slot-controlled airports are often of particular anticompetitive concern in airline mergers. In the most recent major airline merger case in the U.S., the proposed merger between American Airlines and US Airways in 2013, the U.S. Department of Justice expressed significant concern over airline markets associated with DCA. According to its complaint on the AA/US merger case (United States (2013)), the DOJ contends that "passengers to and from the Washington, D.C. area are likely to be particularly hurt. ... Competition at DCA cannot flourish where one airline increasingly controls an essential ingredient to competition. Without slots, other airlines cannot enter or expand the number of flights that they offer on other routes. As a result, D.C. area passengers would likely see higher prices and fewer choices if the merger were approved."

Figure 2: Passenger Trends at Reagan National Airport (DCA)

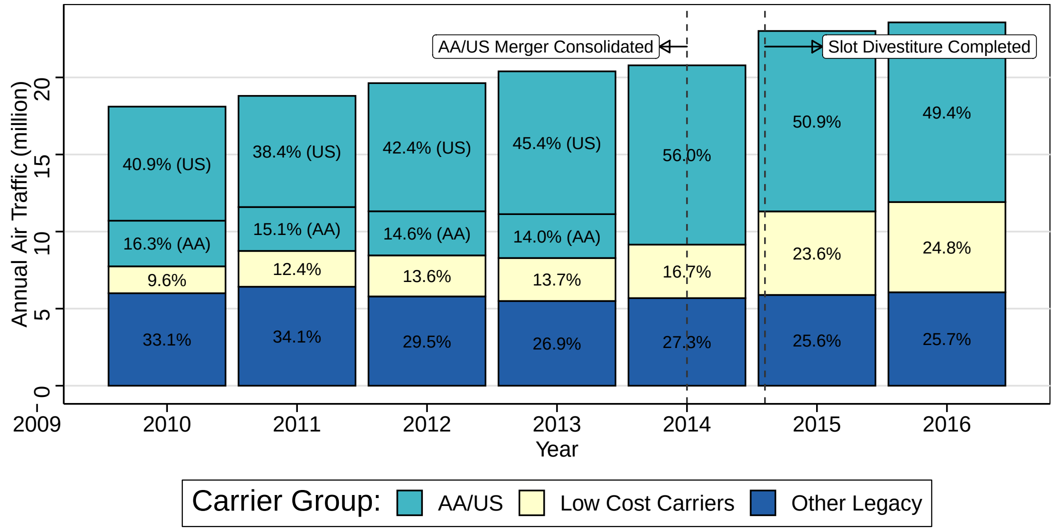

Note: The merger between American and US Airways was completed in December 2013 (marked as a dotted vertical line in the figure). Annual air traffic is the sum of enplaned and deplaned passengers at DCA. The information on annual air traffic is from the Metropolitan Washington Airports Authority (MWAA) website.

Slot divestitures were thus motivated by the desire to alleviate the anticompetitive concern regarding increased market power at this slot-controlled airport. The DOJ and the merging party reached a settlement that required the divestiture of slots and gates to LCCs at seven strategic airports. In this merger divestitures process, only LCCs were approved by the government to purchase the divested assets, based on the government's view that LCCs are more suitable and effective competitors than legacy carriers.[6] The merging firms divested 104 landing slots at DCA, representing, on average, 6 slots per hour and roughly 15% of combined pre-merger slot holdings of merging firms. Figure A1 in the Appendix shows that the divested slots at DCA were roughly evenly distributed across time and the gap of the slot holdings between NewAA and other carriers became much closer after the divestitures. As a result, the landscape of the passenger traffic changed sharply post-merger. As shown in Figure 2, US Airways and American Airlines together held a 59.4% passenger share at DCA in 2013 before the merger. After the merger and slot divestitures, the merging firms' passenger share noticeably decreased, while the passenger share held by LCCs increased post-merger.

[6]Delta was one of those expressing concern about the agreement, commenting, "by prescribing a remedy that forecloses network carriers from competing ... the government will distort the market and deny the traveling public the competitive benefits that only network carriers can deliver."

Figure 3: Illustration of Markets and Flight Segments

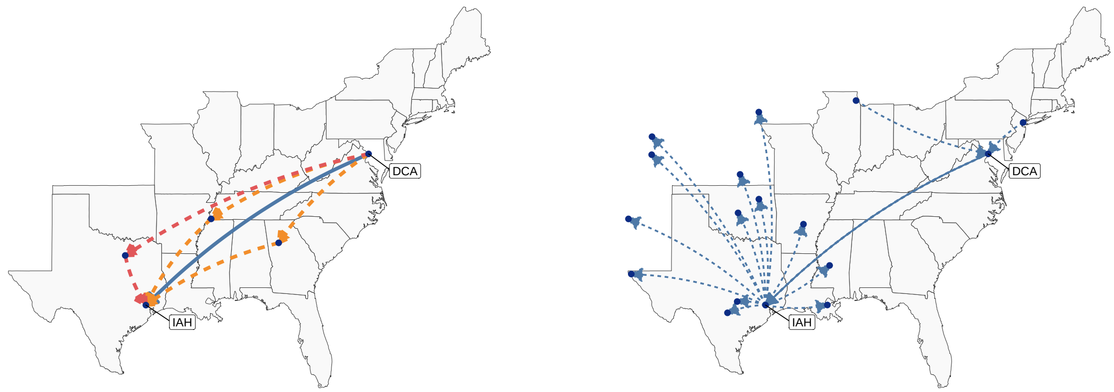

(a) One Market with Multiple Segments (b) One Segment with Multiple Markets

Note: A solid line represents a nonstop product that requires one flight segment, while a dotted line represents a connecting product that requires more than one segments (two segments in this example). An arrow in each line shows a (flight) direction. Different line colors indicate different carriers. While the left panel describes that multiple segments are associated with a single market (DCA⇒IAH), the right panel shows that multiple markets are associated with a single flight segment $(\text{DCA}\rightarrow\text{IAH})_{UA}$.

3 Model

In this section, I develop a two-stage model that captures entry, capacity choice, and price competition at a slot-controlled airport. In the first stage of the model, airline carriers determine a set of flight segments to serve. Then, in the second stage, the carriers allocate their scarce airport slots across those segments and choose product prices. Similar to previous capacity-constrained models (e.g. Besanko and Doraszelski (2004) and Snider (2009)), I allow slot allocation and product pricing to affect marginal cost through a segment's load factor, which enables cross-market interactions in the model.

3.1 Environment

I define an airline market based on potential *directional* trips between two endpoints (e.g, airports or cities) regardless of the number of connections made *en route*, where markets are indexed by $m \in \mathcal{M}$. I also define a flight segment of carrier f, indexed by $s \in \mathcal{S}_f$, as a *directional* nonstop trip of carrier f between two endpoints. While a market is a place where passenger demand takes place, a flight segment for a carrier is an operational unit where the carrier allocates its capacity. I denote $A \Rightarrow B$ as the market of endpoint A to endpoint B, and let $(A \rightarrow B)_f$ be the flight segment from A to B operated by carrier f. Note that the baseline endpoints in this paper are airports unless otherwise stated.

A single market can be associated with multiple flight segments, and vice versa. The map in the left panel of Figure 3 shows multiple products and segments in the DCA to Houston market, DCA⇒IAH. In this example, there are four products—one nonstop and three connecting products operated by three different carriers—and seven different flight segments associated with the DCA⇒IAH market. The map in the right panel of the figure, on the other hand, shows a subset of the flight products operated by United Airlines (UA) that share the flight segment $(\text{DCA}\rightarrow\text{IAH})_{UA}$. In that single segment (United's nonstop flight from DCA to IAH), there are different types of passengers in different markets: 1) those who simply took the flight, 2) those who took prior flights, and 3) those who are traveling on to other destinations. The feature that passengers from different markets can share the same flight segment will be the key to capturing interactions among markets, and I explain how my model incorporates this feature in the subsequent subsection.

3.2 Demand

A nested logit model is used to describe market demand. In market m, a consumer can choose the option of all possible flight products in the market or the option not to fly (outside option). The indirect utility of customer i purchasing flight product j in market m at time t is given by:

$$u_{ijmt} = \underbrace{x_{jmt}\beta - \alpha p_{jmt} + \xi_{jmt}}_{\text{product quality } (\delta_{jmt})} + \zeta_{ijmt}(\sigma) + (1-\sigma)\epsilon_{ijmt}, \tag{1}$$

where x_{jmt} is a vector of observable product characteristics, p_{jmt} is the airfare of product j, and ξ_{jmt} is a product-specific unobservable characteristic. I express $x_{jmt}\beta - \alpha p_{jmt} + \xi_{jmt}$ as a product quality δ_{jmt}. To control for variations in the consumers' tastes across carriers, time and markets, I include carrier, time and *non-directional* market fixed effects (ρ_j, τ_t, and ψ_m, respectively)[7] such that:

$$\Delta\xi_{jmt} = \xi_{jmt} - \rho_j - \tau_t - \psi_m. \tag{2}$$

Controlling for the fixed effects, I assume that $\Delta\xi_{jmt}$ has an i.i.d. distribution. ϵ_{ijmt} denotes consumer i's idiosyncratic preferences for product j, and it follows a type-1 extreme value distribution. Air travel products are nested in one group, and the no-flying option is nested in the other group. The parameter σ governs the degree of substitution between the two groups. As σ goes to one, air travel products become closer substitutes; as σ goes to zero, there are no distinguishable substitution patterns among air travel options, and the nested logit model becomes similar to a simple multinomial logit model. I

[7] For non-directional market fixed effects, I group products in both the A⇒B market and B⇒A market in the same category.

estimate the vector of demand parameters $\boldsymbol{\theta}^{\mathbf{D}} = (\beta', \alpha, \sigma)$.

Let s_{jmt} be the choice probability of product j among alternative products in market m at time t. Then, the probability can be expressed as the following formula under the nested logit model:

$$s_{jmt} = \frac{e^{\frac{\delta_{jmt}}{1-\sigma}}}{D_{mt}} \frac{D_{mt}^{1-\sigma}}{1 + D_{mt}^{1-\sigma}}, \tag{3}$$

where $D_{mt} = \sum_{k \in \mathcal{J}_{mt}} e^{\frac{\delta_{kmt}}{1-\sigma}}$ and $\mathcal{J}_{mt}$ is the set of all products in market m at time t. Finally, denote q_{jmt} as the model-predicted number of passengers who use product j in market m at time t. This equals the product's choice probability multiplied by market size, $M_{mt} s_{jmt}$. Market size is defined as the geometric mean populations (metropolitan statistical areas) of the two endpoints of the market.

3.3 Supply

To describe the optimal behavior of carriers at a slot-controlled airport, I consider markets that are restricted to those originating, ending, or connecting at a slot-controlled airport (DCA in this study). Denote $\mathcal{S}_f$ as the set of flight segments provided by carrier $f \in \mathcal{F}$ to/from the slot-controlled airport. Denote $\mathcal{J}_f(\mathcal{S}_f)$ as carrier f's products that use any segment in $\mathcal{S}_f$. In the model, I assume that product $j \in \mathcal{J}_f(\mathcal{S}_f)$ consists of up to two flight segments—while nonstop flight products require one segment, connecting flights require two.

$$|\mathcal{S}_{fj}| \leq 2, \qquad \forall j \in \mathcal{J}_f(\mathcal{S}_f), \forall f \in \mathcal{F}, \tag{4}$$

where $\mathcal{S}_{fj}$ is the set of flight segments that product j uses.

Capacity Choice Variable: Airport Slot

In this model, the capacity (or the number of available seats) that carrier f provides on a flight segment s at time t is defined as $z_{fst} K_{fst}$—the segment-specific average airplane size z_{fst} multiplied by the average number of daily slots assigned to flight segment s, K_{fst}. As one slot is equivalent to one flight, K_{fst} can be also interpreted as the average number of daily flights on s. I assume that the airplane size z_{fst} is exogenously given from the data, which will be discussed below in detail. On the other hand, carriers are able to allocate K_{fst} endogenously across flight segments in this model. I make several assumptions on K_{fst}, mainly due to model tractability and data limitation.

I assume that K_{fst} is continuous. Allowing K_{fst} to be continuous helps us to quickly solve the optimal solution via first order conditions, while reasonably approximating the carriers' actual slot allocations.[8] In addition, since K_{fst} is interpreted as the average number of daily flights on s, the number needs not be an integer. For example, if a carrier allocates its landing slots on s three times per day on weekdays and two times per day on weekends, then K_{fst} will be $3 \times \frac{5}{7} + 2 \times \frac{2}{7} = 2.714$. Different airline schedules on weekdays and weekends are a very common practice in the airline industry.

Additionally, I make other assumptions on K_{fst}. First, K_{fst} is not time specific. In practice, airport landing slots are time specific (e.g. 7AM and 7PM slots are different). While one might argue that slots at certain hours would be more valuable than other hours, the model does not differentiate slots by time because time specific product prices and capacities are not available in the data. When it comes to DCA, however, slot holdings and flight operations are uniformly distributed across different hours, except late nights and weekend mornings in Figure A3, and daily departure options of a product tend to be spread out across time of day. Second, I assume that an airport slot is a single type, and the model does not differentiate commuter slots from general slots. Instead, the size restriction on commuter slots will be mostly governed by the segment-specific airplane size z_{fst} from the data. Last, flight segments beyond the perimeter rule that require specially exempted slots are not considered in this analysis.

Second Stage: Slot Allocation and Product Price Choices

In the second stage, carrier f, which enters a set of flight segments $\mathcal{S}_f$, allocates its endowed slots $\overline{K_f}$ to the segments $s \in \mathcal{S}_f$ and sets product prices to maximize the carrier's aggregate variable profit, which is given by:

$$VP_f(\mathcal{S}_f, \mathcal{S}_{-f}) = \sum_{\{j,s|j \in \mathcal{J}_f(\mathcal{S}_f), s \in \mathcal{S}_{fj}\}} (p_{jm} - c_{jm}) q_{jm}, \tag{5}$$

$$\sum_{s \in \mathcal{S}_f} K_{fs} \leq \overline{K_f}, \qquad \forall f \in \mathcal{F}, \tag{6}$$

where the time subscript t is suppressed for simplicity. $\mathcal{S}_{-f}$ is a set of flight segments offered by carriers other than f, and c_{jm} is the marginal cost of product j in market m. K_{fs} is defined as the average number of slots per day assigned to flight segment s by carrier f. The inequality condition (6) is a slot constraint that carrier f faces—the sum of slots to be allocated should not exceed the endowed slots of carrier f.

[8] There might be a potential concern on corner solutions for those segments with infrequent daily flights. In the data, however, Figure A2 in the Appendix shows that flight segments that have on average less than 0.5 daily flights (one flight every other day) are very rare (1.2%).

The marginal cost is flight segment specific. For example, a connecting flight consisting of two flight segments incurs costs from both segments. The marginal cost of a product is specified as:

$$c_{jm} = \underbrace{\sum_{s \in \mathcal{S}_{fj}} \left(x_{fs}\gamma_1 + \gamma_2 \Big(\frac{Q_{fs}}{z_{fs}K_{fs}} \Big)^{\nu} \right)}_{\text{flight segment specific}} + \omega_{jm}, \tag{7}$$

where $\mathcal{S}_{fj}$ is the set of flight segments that product j uses, and x_{fs} is a vector of observable characteristics of flight segment s offered by f such as segment distance, hub dummy, or a slot constraint airport dummy. In this specification, assigning a slot on a segment affects the marginal cost via a load factor, $\frac{Q_{fs}}{z_{fs}K_{fs}}$, the ratio of passengers to total available seats on s. While the denominator, $z_{fs}K_{fs}$, is the total number of available seats on s operated by f, the numerator, Q_{fs}, is the total number of passengers flying on segment s offered by f, and is computed as the sum of all passengers that use the products offered by f and that share the segment s:

$$Q_{fs} = \sum_{j \in \mathcal{J}_f(\{s\})} q_{jm}. \tag{8}$$

Then, $\frac{Q_{fs}}{z_{fs}K_{fs}}$ captures the extent to which an average flight on s is full. Due to the nonlinear term ν, the marginal cost in this model increases exponentially in the load factor on s. This is associated with overbooking or hassle fees as an airplane becomes more full. Finally, a product-specific unobservable cost shock is denoted ω_{jm}, which captures product-level characteristics that are not observable to econometricians. For this model section, I assume that all the products in $\mathcal{J}_f(\{s\})$ for each segment s in $\mathcal{S}_f$ are available in the data. In practice, however, there may be some products in $\mathcal{J}_f(\{s\})$ that are not available in the data. In Appendix B.2, I explain how the model should be modified when some products are not available in the data.

Second Stage Optimality Conditions and Cross-Market Interaction

Conditional on the flight segment entry decision in the first stage and under the assumption that the slot constraint (6) holds with equality, the necessary equilibrium conditions in the second stage are characterized as the following system of equations:

$$\frac{dVP_f}{dK_{fs}} = \frac{dVP_f}{dK_{fs'}}, \qquad \forall s, s' \in \mathcal{S}_f, \qquad \forall f \in \mathcal{F}, \tag{9}$$

$$\frac{dVP_f}{dp_{jm}} = 0, \qquad \forall j \in \mathcal{J}_f(\mathcal{S}_f), \tag{10}$$

$$\sum_{s \in \mathcal{S}_f} K_{fs} = \overline{K_f}, \qquad \forall f \in \mathcal{F}. \tag{11}$$

The conditions above are derived from the Lagrangian method applied to the optimization problem in (5) focusing on interior solutions. The optimality condition for slot (9) implies that, conditional on prices, a carrier is able to obtain higher profits by transferring landing slots from abundantly allocated segments to relatively lower allocated segments using up all of the firm's endowed slots. Also, note that while there is no guarantee that a solution is unique, multiple solutions have not been found when starting with different sets of initial conditions.

The optimality condition for price (10) can be written as:

$$\frac{dVP_f}{dp_{jm}} = \underbrace{q_{jm} + \sum_{k \in \mathcal{J}_{mf}} (p_{km} - c_{km}) \frac{\partial q_{km}}{\partial p_{jm}}}_{\Delta VP,\ \text{direct impact of market competition}} \underbrace{- \sum_{l \in \mathcal{J}_f(\mathcal{S}_f)} \frac{\partial c_{lm'}}{\partial p_{jm}} q_{lm'}}_{\Delta VP,\ \text{indirect impact}} = 0, \tag{12}$$

where $\mathcal{J}_{mf}$ is the set of products offered by f in market m and the carrier can offer more than one product in the market (e.g. a nonstop and a connecting flight). The first order condition (12) indicates that a product price change affects the variable profit of the carrier through two channels. First, when the price of product j in market m, p_{jm}, decreases, the variable profit of f changes through within-market competition and its sign will depend on the magnitude of the price elasticities of demand. This channel is widely seen in the literature in the context of differentiated product markets—see Nevo (2001); Villas-Boas (2007); Berry and Jia (2010). If there were no indirect impact term in (12), the equation would be the same as a classical mark-up equation.

In the second channel, changes in load factors affect the variable profit. Specifically, if a decrease in p_{jm} increases the number of passengers, q_{jm}, then the load factors of the flight segments that product j uses increase, and therefore the marginal costs associated with these segments also increase. As there are multiple products in different markets that share the flight segments, all the products in the network that use the same flight segment become relatively more expensive, and the carrier's variable profit thereby futher decreases.[9] This spillover appears in the slot allocation as well in (9). When carrier f transfers some slots from one segment to another, the load factors of those segments will be altered; hence, the marginal costs of all the products associated with those segments and the variable profit of the carrier will also be altered. Note that a change in price (or slot) simultaneously alters both the first order conditions for price and slots, as the optimality conditions (9) to (10) are interconnected. In equilibrium, we find the optimal slot allocation and product prices that satisfy those optimality conditions.

One remark is that the model can be potentially extended in such way that product demand is affected by slot allocation. In the current model setup, allocating slots on a segment only affects the supply side via a load factor. As suggested by Berry and Jia (2010), however, consumers may prefer a product with more daily departure options. Since one additional slot on a flight segment proportionally increases the number of daily departures of multiple products, it is natural to extend the model for the demand side. For this, I explore the demand specification when daily departure options are added to the baseline case. In addition, I describe the extension of the model and the challenges that it creates in Appendix D.2.

Under the marginal cost functional form (7), equation (12) can be written as a simple matrix notation for market m:

$$\mathbf{\Omega_m^{-1} q_m} + \mathbf{p_m} = \mathbf{c_m} + \frac{\mathbf{dc_m}}{\mathbf{dQ}}, \qquad \forall m \in \mathcal{M}. \tag{13}$$

where the bold face font in the equation indicates that variables are vector or matrix. $\mathbf{\Omega_m}$ is the element-wise multiplication of the response matrix containing the derivatives of market shares with respect to price, and the matrix indicating whether products i and j are owned by the same carrier. The (i,j) element of Ω_{ijm} is given by:

$$\Omega_{ijm} = \begin{cases} \frac{\partial q_{im}}{\partial p_{jm}} & \text{if } i,j \in \mathcal{J}_{mf} \\ 0 & \text{otherwise,} \end{cases} \tag{14}$$

and $\frac{\mathbf{dc_m}}{\mathbf{dQ}}$ is a vector with jth element

$$\Big[\frac{\mathbf{dc_m}}{\mathbf{dQ}}\Big]_j = \sum_{s \in \mathcal{S}_{fj}} \gamma_2 \nu \Big(\frac{Q_{fs}}{z_{fs} K_{fs}}\Big)^{\nu}. \tag{15}$$

Equation (13), an extended version of the widely used standard pricing equation in a differentiated product market context, is used to estimate the vector of marginal cost parameters $\boldsymbol{\theta}^{\mathbf{C}} = (\gamma_1', \gamma_2, \nu)$. In Appendix B.1, I provide more detailed technical information

[9] In the data, 5.8 products share the same flight segment at DCA, on average.

on how to derive (13) from (12).

First Stage: Flight Segment Choices

In the first stage, carriers choose the set of flight segments that yield the highest expected profits among the alternatives. Following Eizenberg (2014) and Wollmann (2018), I assume that carriers only know the distributions of demand and marginal cost shocks (ξ and ω, respectively), while they know the realized unobservable shocks for fixed costs in this stage. This implies that carriers know their variable profit at a slot-controlled airport in expectation, while they exactly know the incurred fixed costs. Given others' segment offerings $\mathcal{S}_{-f}$, the expected profit of carrier f that offers flight segments $\mathcal{S}_f$ can be expressed as:

$$\Pi_f(\mathcal{S}_f, \mathcal{S}_{-f}) = \mathbb{E}_{\xi,\omega}\left[VP_f(\mathcal{S}_f, \mathcal{S}_{-f})\right] - \sum_{s \in \mathcal{S}_f} F_{fs}, \tag{16}$$

where F_{fs} is the fixed cost that carrier f pays every quarter to operate a nonzero frequency of flight services on segment s. I assume that the fixed cost has the following form:

$$F_{fs} = X_{fs}\eta + \phi_{fs}, \tag{17}$$

where X_{fs} is a vector of carrier-segment specific characteristics that affects fixed costs and ϕ_{fs} is a fixed cost error term that is unconditionally mean zero. In particular, X_{fs} includes a constant, LCC dummy (LCC_f), the ratio of f's total endowed slots to the total endowed slots by any carrier at DCA ($SlotRatio_f$), and the airport presence of carrier f on the non-DCA endpoint of segment s ($Presence_{fs}$). I include $SlotRatio_f$, as carriers with a large number of available slots at the slot-controlled airport may have a cost advantage of initiating a new service by more flexibly adjusting flight schedules.

The fixed cost parameter vector, $\boldsymbol{\theta}^{\mathbf{F}} = \eta'$, is estimated via a partial identification approach. To estimate the parameters, I use the revealed preference to form moment inequality conditions by adding or removing a flight segment to/from a carrier's network. The literature recognizes that the information assumption imposed in this model generates a selection problem (Holmes (2011); Eizenberg (2014); Wollmann (2018)). In brief, carriers may self-select into flight segments that have relatively cheaper fixed costs that are unobservable to econometricians. In Section 5.3, I will discuss how I address the selection issue when estimating the fixed cost parameters.

4 Data

This section discusses the data used in this study and the sample selection procedure. I combined the two publicly available datasets from the Department of Transportation – 1) the Airline Origin and Destination Survey (DB1B), and 2) the Air Carrier Statistics (T-100). While I use the entire sample for demand and marginal cost parameter estimations, a subset of the products associated with DCA is used for the fixed cost estimation and counterfactual analysis.

4.1 Airline Origin and Destination Survey (DB1B)

DB1B contains a 10% sample of all air travel passenger itineraries in the U.S. domestic airline industry. It includes the following information: the origin, destination, and connecting (if any) airports; the number of passengers; ticket prices; and the ticketing and operating carriers on each itinerary. To capture the products and markets before the AA/US merger was consummated in December 2013, I use a dataset spanning from the first quarter of 2012 to the last quarter of 2013. The dataset comprises 7,073 markets (across time) taken from the set of routes connecting the largest 100 U.S. domestic airports based on the number of passenger boardings in 2012.

I drop the itineraries that have more than one connection, as less than 1% of passengers use itineraries with more than one stops in the data. I also drop itineraries for which the ticket prices are outside the range of $12.50 to $1,250, as these are likely coding errors. I aggregate the tickets with the same itineraries and take the passenger-weighted average prices as a representative price for a product. Similar to Ciliberto and Williams (2014), I further drop products with fewer than 200 passengers in a market, as those are not likely to be effective competing services in the market. Note that I treat products that have the same origin and destination but different connections operated by the same carrier as distinct products. For example, the connecting products (DCA→Atlanta (ATL)→IAH)$_f$ and (DCA→Dallas (DFW)→IAH)$_f$ are distinct, while they are in the same market, DCA⇒IAH.

Panel A of Table 1 reports the descriptive statistics of the product characteristics from DB1B. In column "All Products," where I use the entire sample, we see that the average ticket price is $246 and that on average 1,749 passengers fly with the same product. While 11.5% of products are nonstop, nonstop passengers consist of more than 75% of the total passengers. "Daily Departures" is defined based on the average number of daily flights of a product, and I construct the variable following the procedure suggested by

Table 1: Descriptive Statistics

	All Products		DCA Products	
Variable	Mean	Std. Dev.	Mean	Std. Dev.
Ticket Price ($100)	2.463	0.643	2.475	0.684
Passengers (1,000)	1.749	5.818	1.791	6.096
Nonstop (Dummy)	0.115	0.319	0.108	0.310
Daily Departures	3.388	2.532	4.055	2.702
Distance (1,000 miles)	1.322	0.651	1.173	0.674
Presence (Origin)	0.298	0.269	0.242	0.247
Slot Controlled (Dummy)	0.143	0.350	1.000	0.000
Market Size (Million)	3.005	2.262	3.443	2.171

(a) Panel A. Product Level (DB1B)

	All Segments		DCA Segments	
Variable	Mean	Std. Dev.	Mean	Std. Dev.
Quarterly Departures	279.230	250.751	309.461	277.525
Available Seats (1,000)	31.956	34.934	29.956	36.014
Passengers (1,000)	25.824	28.867	22.224	27.395
Airplane Size	114.641	42.322	96.512	41.771
Load Factor	0.803	0.108	0.741	0.124
Segment Distance	0.915	0.593	0.772	0.578

(b) Panel B. Flight Segment Level (T-100)

The descriptive statistics in Panel A for all products and DCA products are based on 358,880 and 14,819 observations, respectively. Similarly, the numbers of unique observations of Panel B for all products and DCA products are 42,958 and 1,650 respectively. The markets are based on airport pairs.

Berry and Jia (2010).[10] An average product has 3.38 daily flights. Presence at an origin airport, “Presence (Origin),” is defined by the ratio of the number of routes that a carrier serves nonstop to the total number of routes that any carrier serves nonstop at the origin airport. Of the products considered, 14.3% depart from, arrive at, or transfer at a slot-controlled airport. I restrict the data to those products associated with DCA, and the summary statistics of the subsample are reported in the column “DCA Products.” While it seems that there is no distinguishable difference in prices, the number of passengers, and the fraction of nonstop products between “All Products” and “DCA Products”, “DCA Products” tend to have more daily departure options and less total miles flown than the average product in the industry.

[10]Due to the lack of data, the Daily Departure values for 15.2% of connecting products are not extracted by using the method suggested by Berry and Jia (2010). They are imputed as the geometric mean of the daily departures of the two segments of a connecting product.

4.2 Air Carrier Statistics (T-100)

The T-100 dataset provides monthly flight-level information, including the number of passengers, available seats, types of aircraft, and flight frequency on a segment. To match the units of T-100 and DB1B, different aircraft types of the same carrier within a segment are aggregated at the quarterly level. Regional affiliates shown in T-100 are converted to their affiliated ticketing carriers by matching the operating carriers shown in DB1B.

Panel B of Table 1 summarizes the flight segment characteristics. An average flight segment has 31,956 seats, approximately 80% of which are filled with passengers, and the average airplane size is 114.64 (seats). In terms of the DCA segment, the average airplane size, load factor, and segment distance are smaller than the industry average. This is partially due to the fact that approximately 20% of slots at the airport are designated as commuter slots, and partially due to the perimeter rule, which restricts flight operations to within 1,250 miles. Small-sized aircraft are more appropriate for short-haul routes where flights tend to be less full.

4.3 DCA Products

While I use the sample "All Products" to estimate the demand and marginal cost parameters, I further refine the "DCA Products" shown in Table 1 to conduct the counterfactual analysis as well as to estimate the fixed cost parameters. I start with the set of all products that include any DCA segments within 1,250 miles in the second quarter of 2013, and I denote $\mathcal{M}$ as the set of markets to which those products belong. In this way, those long haul DCA segments beyond the perimeter rule using the specially exempted slots are excluded.

Some markets in $\mathcal{M}$ may not be to/from DCA, and some products of those markets may not be in the "DCA Products." For example, suppose there are two products in ATL⇒Boston (BOS) market—1) a connecting product $(\text{ATL}\rightarrow\text{DCA}\rightarrow\text{BOS})_f$ by carrier f, and 2) a nonstop product $(\text{ATL}\rightarrow\text{BOS})_{f'}$ by carrier f'. Because the first connecting product uses DCA segments, the ATL⇒BOS market should be included in $\mathcal{M}$. However, the nonstop product is not included in the "DCA Products," as it does not contain DCA flight segments. While adding more products to the sample may allow a more realistic counterfactual analysis, doing so would also increase the computational burden of finding the optimal slot allocation and product prices. For a viable counterfactual exercise, I reduce the number of products by taking the simplifications described below.

First, any product of markets in $\mathcal{M}$ that does not contain any DCA segment is not included in "DCA Products" (e.g. the nonstop flight in ATL⇒BOS in the example above).

Table 3: Statistics of Refined DCA Products in 2013Q2

	Products		Passengers		Segments/Slots	
Carrier	N.	Nonstop (%)	Pass. N. (mil.)	Nonstop (%)	Seg. N.	Endowed Slot
US Airways	315	31.11	1.38	82.48	98	369.59
American	84	16.67	0.48	76.95	14	91.55
Delta	195	7.18	0.49	65.40	14	88.56
United	105	11.43	0.33	78.41	12	62.89
JetBlue	23	34.78	0.25	94.04	8	32.87
Southwest	71	14.08	0.20	67.69	10	27.08
Total	793		3.12		156	672.54

Note: In the sample, there are 28 composite products in four composite markets, accounting for 7% of DCA passenger traffic.

Instead, its product quality (δ), defined in (1), is calculated from a demand estimation, and is used to adjust the consumer choice probability of "DCA Products" in (3). In the earlier example of the ATL⇒BOS market, the choice probability of the connecting product $(\text{ATL}\rightarrow\text{DCA}\rightarrow\text{BOS})_f$ in (3) is calculated as if there is the nonstop product by taking into account the quality of the nonstop product. In this way, market competition can be reasonably described, while we remove a subset of products from the sample.

Second, I create composite products for connecting services of a carrier that share the same DCA segment. There are a number of connecting flights for which one end is DCA and the other end is located in the Midwest or West, connecting at large hub airports (e.g. the itinerary DCA→Detroit (DTW)→Los Angeles (LAX) offered by American Airlines). While those connecting products have a negligible presence in their corresponding markets, the sum of those passengers has non-negligible effects on the load factors on its DCA segment (7% of the total passengers using any DCA segments fall into this group). Therefore, instead of eliminating those products from the sample, I combine them into a single composite product. To be specific, I combine those connecting products that share the same flight segment ending at DCA and have a market (nonstop) distance beyond the perimeter rule threshold. Figure A4 illustrates two examples of composite products—one is DL's composite product connecting at Atlanta, and the other is AA's composite product connecting at Dallas.[11] I assume that those composite products are in one of two composite markets depending on the location of connecting airports (south or north), which creates a moderate market competition.

Table 3 reports the number of products, passengers, and segments of those in the "DCA Products" category by carrier in the second quarter of 2013. The subsample contains 793 products, which represent 3.12 million passengers, and 79% of those passengers

[11] Note that a connecting flight in the DCA to Miami (MIA) market is not aggregated, as the market distance is less than 1,250 miles.

use nonstop products. On average 5.8 products share the same flight segment. I calculate a carrier-level slot endowment at DCA by counting the total number of quarterly departures to/from DCA segments shown in T-100 for each carrier, excluding the flight segments beyond the perimeter rule threshold. Then, I divide this number by 91 days so that its unit becomes the average number of slot holdings per day. The result is shown in the last column. There are 156 carrier-specific flight segments in this sample, and US Airways and American held 54.9% and 13.6% of the slots pre-merger, respectively. The calculated slot holdings are consistent with the DOJ's complaint claiming that "US Airways holds 55% of slots at DCA pre-merger, and the proportion would increase to 69% when AA/US merge."[12]

5 Estimation

This section discusses the estimation of the model parameters. While the demand and marginal cost parameters are point estimated, the bounds on the fixed cost parameters are estimated by using moment inequalities.

5.1 Demand

As proposed by Berry (1994), the estimation of the demand parameters $\boldsymbol{\theta}^{\mathbf{D}} = (\beta', \alpha, \sigma)$ is based on the regression equation (18) inverted out from observed market shares in (3):

$$\log(s_{jmt}) - \log(s_{0mt}) = x_{jmt}\beta - \alpha p_{jmt} + \sigma \log(s_{Gjmt}) + \xi_{jmt}, \tag{18}$$

where s_{0mt} is the market share of the not flying option, and s_{Gjmt} is the within-group market share of product j. Given the model's information assumption that unobserved demand and cost shocks are not realized when carriers make flight segment entry decisions, selection is not present unless carriers regret their entry decision and reverse their entry choice.

There are two parameters for which the estimates are subject to endogeneity bias in the equation above—the price coefficient (α) and the nesting parameter (σ). Since carriers may account for unobservable product characteristics (ξ_{jmt}) when they make price

[12] The validity of the slot holdings measure that I construct can be checked using an alternative source. Since June 2017, the FAA has uploaded information on slot holdings at slot-controlled airports (Figure A3 is based on this information). By considering the actual slot divestitures, I backed out the number of slots that each carrier would have held before the merger. While the backed-out number is slightly greater than what I construct, as it includes slots that are exempted from the perimeter rule, I find that the two measures are consistent.

decisions, prices are likely endogenous, and s_{Gjmt} is mechanically correlated with ξ_{jmt}. I address the endogeneity concerns with rich fixed effects specifications and instrumental variables (IV).

First, I include carrier, quarter, and *non-directional* market fixed effects to control for unobservables that are constant along those dimensions. Second, I use instrumental variables that are correlated with prices but uncorrelated with the unobservable product characteristics (ξ_{jmt}) in (18). I exploit market structure changes as a shock to shift markups, such as the number of within market nonstop carriers and whether any LCC exists within market. The validity of these instruments is based on the timing assumption of the model. Since a flight segment entry decision occurs before unobservable demand shocks are realized, market structure changes via entry/exit are not correlated with those unobservables. Additionally, I use LCC dummy variables interacted with slot-controlled airport dummies as a cost shifter IV, since LCCs at slot-controlled airports are faced with a tight capacity constraint leading to an increase in operational cost. These demand estimates allow me to obtain demand residuals $\hat{\xi}_{jmt}$, and the distribution of the residuals will be used to calculate the expected variable profit for the fixed cost estimation and counterfactual analysis.

5.2 Marginal Cost

The estimation of the marginal cost parameters $\boldsymbol{\theta}^{\mathbf{C}} = (\gamma_1', \gamma_2, \nu)$ is based on the pricing equation (13). The j'th element of the equation can be written as:[13]

$$\begin{aligned}\left[\boldsymbol{\Omega}_{\mathbf{m}}^{-1}\mathbf{q}_{\mathbf{m}} + \mathbf{p}_{\mathbf{m}}\right]_j &= \left[\mathbf{c}_{\mathbf{m}} + \frac{\mathbf{dc}_{\mathbf{m}}}{\mathbf{dQ}}\right]_j \\ &= \sum_{s \in \mathcal{S}_{fj}} \left(x_{fs}\gamma_1 + \gamma_2(1+\nu)\Big(\frac{Q_{fs}}{z_{fs}K_{fs}}\Big)^{\nu} \right) + \omega_{jm}. \end{aligned} \tag{19}$$

I adapt a nonlinear IV-GMM method to estimate both the linear and nonlinear parameters in (19). Similar to the BLP demand estimation proposed by Berry et al. (1995), the nonlinear IV-GMM has inner and outer loops. In the outer loop, I search over the nonlinear parameter ν that minimizes the objective function of the generalized method of moments (GMM):

$$\hat{\nu} = \arg\min_{\nu} (Z'\omega) W^{-1} (Z'\omega)', \tag{20}$$

[13] The equation will be slightly different if missing products exist. In Appendix B.2, I explain this in detail.

where Z and ω are the stacked versions of the set of instrumental variables and the marginal cost error term, respectively, and W^{-1} is the inverse of the weighting matrix. Conditional on ν, in the inner loop, γ_1 and γ_2 are recovered by an IV regression. Then, the marginal cost residuals $\widehat{\omega}$ are fed into the moment condition (20) until we find the $\widehat{\nu}$ that minimizes the GMM objective function.

The load factor term is endogenous in (19). Specifically, a lower unobservable marginal cost tends to cause a lower price and a higher load factor holding slot allocation fixed. To address the endogeneity problem, in the spirit of Fan (2013), I introduce an IV that measures the market competitiveness of the "neighboring markets"—the market of the products that share the same flight segment. To illustrate this, consider two products in the right panel of Figure 3 that share the flight segment $(\text{DCA}\rightarrow\text{IAH})_f$—one is a nonstop flight in the DCA⇒IAH market, and the other is a connecting flight in the DCA⇒Denver (DEN) market connecting at IAH. When the competition in the DCA⇒DEN market changes, the number of passengers using the connecting flight changes and, hence, the load factor of $(\text{DCA}\rightarrow\text{IAH})_f$ affects both products. The change in load factor for the nonstop product in $(\text{DCA}\rightarrow\text{IAH})_f$ is not caused by marginal cost unobservables of the product but by the market structure changes in DCA⇒DEN market that the connecting flight belongs to. Given the information assumption that the flight segment entry decision of carriers is chosen prior to the realization of marginal product shocks, ω, the exclusion restriction holds for the introduced instrumental variable.

The variables related to market competitiveness in the "neighboring market" are constructed in the following way. When product j is a nonstop product, I list the set of products other than j that share the flight segment that j uses. Then, I find the market of the product that brings the highest number of passengers to the flight segment other than j, which I call a "neighboring market". Similar to the instrumental variables in the demand estimation, I use the number of nonstop carriers and any LCC dummy in the "neighboring market" as IVs for the endogenous load factor variable of j. In terms of a connecting product, analogously, I separately calculate the IVs of the two flight segments that the product has.

5.3 Fixed Cost

Given the demand and marginal cost point estimates $\hat{\theta}^{\mathbf{D}}$ and $\hat{\theta}^{\mathbf{C}}$, the estimation of the fixed cost parameters $\theta^{\mathbf{F}} = \eta'$ relies on a partial identification strategy. Due to the discrete nature of an entry decision, there is no guarantee of a unique equilibrium for choosing the set of flight segments—see Eizenberg (2014). Therefore, it is challenging to obtain the point estimates of the fixed cost parameters without an additional assumption (e.g. the sequential order of carriers for entry decisions). Instead, following a growing literature on fixed cost estimations (Eizenberg (2014); Ho and Rosen (2015); Wollmann (2018)), I adopt a partial identification approach, and use the "DCA Products" sample discussed in Section 4.3.

I assume that the observed set of flight segments constitutes a pure strategy Nash equilibrium and that any unilateral deviation from the set of flight segments will not increase the expected profits of a carrier. To illustrate this, suppose that $\mathcal{S}_f^*$ and $\mathcal{S}_{-f}^*$ are the observed sets of flight segments offered by f and carriers other than f, respectively, at a slot-controlled airport. Removing a flight segment s from the existing airline network does not increase f's expected profit, and this condition forms an upper bound on the fixed cost F_{fs}, $\overline{F}_{fs}$:

$$F_{fs} \leq \mathbb{E}_{\xi,\omega}\Big[VP_f(\mathcal{S}_f^*, \mathcal{S}_{-f}^*) - VP_f(\mathcal{S}_f^* \backslash \{s\}, \mathcal{S}_{-f}^*)\Big] \equiv \overline{F}_{fs}, \qquad \forall s \in \mathcal{S}_f^*, \forall f \in F, \tag{21}$$

where $\mathcal{S}_f^* \backslash \{s\}$ is the set of flight segments offered by f excluding the segment s. Analogously, a lower bound, $\underline{F}_{fs}$, on the fixed cost can be formed by using the condition that adding a new flight segment s to the existing network cannot make f better off:

$$F_{fs} \geq \mathbb{E}_{\xi,\omega}\Big[VP_f(\mathcal{S}_f^* \cup \{s\}, \mathcal{S}_{-f}^*) - VP_f(\mathcal{S}_f^*, \mathcal{S}_{-f}^*)\Big] \equiv \underline{F}_{fs}, \qquad \forall s \in \mathcal{S}_f \text{ and } s \notin \mathcal{S}_f^*, \forall f \in F, \tag{22}$$

where $\mathcal{S}_f^* \cup \{s\}$ is the set of flight segments of f that adds a new segment s to $\mathcal{S}_f^*$. Although a wide range of moment conditions are possibly formed by removing/adding multiple flight segments, I use only one segment deviation to ease the computational burden.

When a carrier adds a new flight segment to its airline network, the carrier pays a fixed cost and commits to providing nonstop service on the segment. There may be a set of potential connecting service offerings when the new flight segment is added, and those service offerings will depend on the characteristics of the origin/destination airports (e.g. hubs), the carrier's network type (e.g. hub-to-spoke or point-to-point system) and the structures of the markets for which the new connecting services are possible.

In this study, however, when carriers add a new segment to their existing network, I assume that they offer only a nonstop product and no connecting products for two reasons. First, since US Airways is the only carrier that considered DCA as a hub pre-merger, those itineraries connecting at DCA are mostly from that carrier—94% of one-stop passengers connecting at DCA were flown by the US Airways. When examining the data, the carrier connected most of its flight segments at DCA, leaving only a few small communities with a small-sized demand as segments to be added. This implies that the connecting passenger flows from those new flight segments based on the small community would be very small. Second, in terms of carriers other than US Airways that do not consider DCA as a hub, it is not likely that they offer products connecting at DCA for transferring passengers on the new segments. The only remaining scenario in which these other carriers could generate connecting passengers is when the non-DCA endpoint of the new segments is their hub airport. However, the data indicate that those carriers already linked their hubs to DCA, and that there are no new segments considered as their hubs.

Airplane Size for a New Segment

While the available seats for an existing segment can be directly obtained from the data (T-100), the airplane size multiplier of a new segment s, z_{fs} in (7) needs to be predicted. For this prediction, I use the sample of those flight segments which daily flight is at least one from 2012Q1 to 2013Q4. Airline carriers allocate their airplanes on a segment based on a number of factors, including aircraft fleet portfolio, flight length, and contracts with regional operators. Typically, a smaller airplane is used for a short-haul distance route. To predict the airplane size, I regress the segment-level average airplane size on distance, and distance squared. In addition, I include fixed effects of both origin-carrier pairs and destination-carrier pairs in the regression to control carrier-airport specific unobservable characteristics that affect fleet choices.

I report the regression result in Table A1 in the Appendix. Airplane size tends to increase as the segment distance increases but in a diminishing way. Airplane size is likely to be smaller at slot-controlled airports, partly because operations at those airports are affected by perimeter rules. Furthermore, airplanes at high-tourism airports (airports in Fort Lauderdale, Orlando, or Las Vegas) or international hubs tend to be large. The preferred specification used to predict airplane size is column (3) where origin-carrier and destination-carrier specific fixed effects denote the specification with the highest predicted power. Last, when the model relaxes the assumption that all products are available in the data, the load factor on s originating from unavailable products also needs to be predicted in order to describe available seats and load factors realistically. In Appendix B.2, I explain the prediction procedure for doing so.

Selection Issue

The information structure of this model setup entails a selection problem. Recall that the fixed cost parameter has a mean-zero error term ϕ_{fs} in (17) that affects the entry decision of flight segment s by carrier f. Since the error term is observed by carrier f but not by econometricians, the carrier may selectively enter those flight segments with lower fixed costs. In other words, ϕ_{fs} is not mean zero conditional on those segments that the carrier decides to enter, while its unconditional mean is zero. This leads to a biased fixed cost parameter estimation.

There are several methods introduced in the literature to address the selection problem. Ciliberto and Tamer (2009) assumes the parametric distribution of fixed cost error terms and obtains a bound on fixed costs by computing the probability of observing equilibrium offerings. However, this method is computationally very expensive in my model setting because it needs to analyze every possible combination of flight segments. Next, Eizenberg (2014) assumes that conservatively wide bound exists for any fixed cost, and he uses the wide bound for counterfactual inequality conditions to obtain unbiased parameter bounds. Following Eizenberg (2014), I address the selection issue by assuming that the support of a fixed cost is the minimum and maximum of the support of the expected change in variable profit coming from a single flight segment change:

Assumption 1 $|F_{fs}| < \infty$ and $\left[\min(F_{fs}), \max(F_{fs})\right] \subset \left[\min(\Delta VP_f), \max(\Delta VP_f)\right]$, $\forall f \in F$ where ΔVP_f is an expected change in variable profit due to the elimination or addition of a single flight segment by carrier f.

Under the Assumption 1, I use $\min(\Delta VP_f)$ and $\max(\Delta VP_f)$ for the lower bound of (21) and the upper bound of (22), respectively. Then, I combine (21) and (22) and take the unconditional expectation for F_{fs}:

$$\mathbb{E}\left[LB_{fs}\right] \le \mathbb{E}\left[F_{fs}\right] \le \mathbb{E}\left[UB_{fs}\right] \tag{23}$$

$$\mathbb{E}\left[LB_{fs}\right] \le X_{fs}\eta \le \mathbb{E}\left[UB_{fs}\right] \qquad (\because \mathbb{E}\left[\phi_{fs}\right] = 0)$$

where

$$LB_{fs} = \begin{cases} \underline{F}_{fs} & \text{if } s \in \mathcal{S}_f \text{ but } s \notin \mathcal{S}_f^* \\ \min(\Delta VP_f) & \text{otherwise,} \end{cases} \qquad UB_{fs} = \begin{cases} \overline{F}_{fs} & \text{if } s \in \mathcal{S}_f^* \\ \max(\Delta VP_f) & \text{otherwise.} \end{cases}$$

Moments and Inferences

Similar to Pakes et al. (2015) and Wollmann (2018), I form the following moment inequality conditions from (23):

$$\frac{1}{N} \sum_{s \in \mathcal{S}_f} \left[X_{fs}\eta - LB_{fs}\right] \ge 0, \qquad \forall f \in \mathcal{F} \tag{24}$$

and

$$\frac{1}{N} \sum_{s \in \mathcal{S}_f} \left[UB_{fs} - X_{fs}\eta\right] \ge 0, \qquad \forall f \in \mathcal{F} \tag{25}$$

where N is the number of new and existing flight segments of f. Following the procedure suggested by Andrews and Soares (2010), the confidence sets of fixed cost parameters are estimated. The procedure is as follows: 1) for a given set of η, compute the objective function $Q_n(\eta) = \sum_i \left(\left[\frac{\bar{m}_i(\eta)}{\hat{\sigma}_i(\eta)} \right]_- \right)^2$, where $\bar{m}_i(\eta)$ and $\hat{\sigma}_i(\eta)$ are the ith moment's sample mean and standard error, and $[x]_-$ operator is defined as x if $x < 0$, and equals 0 otherwise. 2) Draw a large number of bootstrap samples R and compute the objective function for sample at η. 3) Compute a critical value ($c_{1-\alpha}$) at the $(1-\alpha\%)$ quantile of the distribution of the objective based on the bootstrapped samples. 4) Include η in the bound if $Q_n(\eta) < c_{1-\alpha}$. 5) Repeat steps 1-4 for all possible parameter vectors.

6 Estimation Results and Model Fit

6.1 Demand

Table 4 reports the estimation results of the demand system. The first three columns show the baseline demand estimates under different market definitions (airport-airport pair in (1) and (2), and city-city pair in (3)), and the demand parameter estimates are sensible. All other things being equal, air travel passengers strongly prefer nonstop to connecting flights, and carriers with higher airport presence are preferred. As the non-directional market fixed effects are included in the regression, factors constant at the market level such as market distance are not identified in the demand specification. Instead, I add to the regression "Extra-miles," measured by the total distance flown divided by the non-stop market distance.[14] The estimates suggest that connecting flights with longer detours (more extra-miles) are less preferred to consumers, but this effect diminishes as the ratio increases. As expected, airfares negatively affect consumer demand, and nesting parameters are moderately estimated in a way that product substitutions within a nest are more likely than those across nests.

Labeled "Frequency in Demand," the last two columns report the estimation result when daily departure-related variables are added to the baseline specification. Passengers prefer a flight that has more daily departure options, and this preference is much stronger for a nonstop flight than for a connecting flight. For 15.2% of the connecting products in the sample, there is no information on daily departures, and for each of those connecting products, I impute this value as the geometric mean of the daily flight fre-

[14]For example, the total distance flown for a connecting flight from ATL to ORD connecting at DCA, is 1,159 miles—the sum of the two flight segments distances (547 miles for ATL→DCA and 612 miles for DCA→ORD). As the nonstop flight distance between ATL→ORD is 606 miles, the value of Extra-miles for this product is 1.893.

Table 4: Demand Parameter Estimation Results

	(1) Baseline	(2) Baseline	(3) Baseline	(4) Frequency in Demand	(5) Frequency in Demand
Nonstop	0.128***	0.502***	0.367***	0.157***	0.496***
	(0.005)	(0.051)	(0.073)	(0.006)	(0.035)
Presence (Origin)	0.026***	0.159***	0.148***	0.013***	0.105***
	(0.001)	(0.011)	(0.013)	(0.002)	(0.010)
Extra-miles	−0.031***	−0.520***	−0.630***	−0.106***	−0.570***
	(0.011)	(0.093)	(0.162)	(0.012)	(0.077)
Extra-miles squared	0.008**	0.132***	0.160***	0.028***	0.152***
	(0.004)	(0.023)	(0.040)	(0.004)	(0.020)
Fare (α)	−0.058***	−0.339***	−0.352***	−0.071***	−0.384***
	(0.002)	(0.026)	(0.031)	(0.003)	(0.025)
Nesting (σ)	0.955***	0.794***	0.814***	0.932***	0.765***
	(0.001)	(0.021)	(0.033)	(0.002)	(0.018)
Daily Departure (DDep)				0.013***	0.047***
				(0.0005)	(0.003)
DDep X Nonstop				0.016***	0.051***
				(0.001)	(0.006)
DDep X Connecting (Imputed)				−0.006***	−0.026***
				(0.0004)	(0.002)
OLS or IV?	OLS	IV	IV	OLS	IV
O-D Pair (Airport or City)	Airport	Airport	City	Airport	Airport
Median Elasticity	-2.699	-3.535	-4.101	-2.180	-3.514
Diversion Ratio					
(Nonstop → Nonstop)	0.722	0.555	0.560	0.696	0.529
(Nonstop → Connecting)	0.201	0.148	0.191	0.192	0.140
Observations	358,880	358,880	291,930	358,880	358,880

Note: ∗∗∗ < 0.01, ∗∗ < 0.05, ∗ < 0.10. All the specifications include time, carrier, and route fixed effects.

quency of the two segments. I separately estimate the impact on preference by adding its interaction term. While the estimated interaction coefficient is negative, it varies with how the functional form of the imputation is defined. The willingness to pay of a passenger for an additional daily departure of a nonstop (connecting) product is calculated as \$39.0 (\$8.17), in dollar terms. When all flights additionally increase one daily departure, the aggregate demand would increase by 8.8%. This result is relatively smaller than the previous finding in Berry and Jia (2010) that claims that adding one daily departure to all flights increases the aggregate demand by 16% in 2006.

Sensible elasticities and a diversion ratio are calculated under the demand system. The median own-price elasticities range from −4.10 to −3.51 conditional on columns with

IVs. For example, the median and mean elasticities in Ciliberto and Williams (2014) and Li et al. (2019) range from −3 and −4. Compared to an airport-airport pair, demand tends to be more elastic under the city-city pair market definition, as there are a wide range of products that passengers can choose in the city-pair market. For those markets where there are at least two nonstop carriers and one connecting carrier, I report the diversion ratio to explore substitution patterns when a nonstop product price increases.[15] The mean diversion ratio of nonstop to nonstop is three to four times greater than that of nonstop to connecting; this indicates a strong preference for nonstop service. The distribution of the demand residual for the baseline (2) is reported in Figure A5a.

6.2 Marginal Cost

Table 5 reports flight segment-level marginal cost estimates. All the results in the table are based on the airport-pair market definition, and the results based on the city-pair market definition are not included, as it is less straightforward to allocate slots to city-based segments at a slot-controlled airport. While load factors are overestimated when the mark-up values are calculated based on the OLS demand results (in column (1) and (3) in Table 5), the load factors under the IV-based demand estimates are estimated in a way that it moderately affects the marginal cost. For example, the combination of ν and γ_2 in column (2) indicates that a load factor of 0.8 for a flight (sample average) increases the marginal cost by \$15.8, while a load factor of 1.0 increases it by \$24.6. The estimation suggests that the marginal cost increases exponentially as flights become more full. In terms of other variables, marginal cost increases in distance with a small but an increasing rate. A distance of 1,000 miles of a flight is associated with a marginal cost increase of \$26.8. The role of slot constraints and hub airports on marginal cost seems small.

Instrumental variable techniques are used to deal with the endogeneity of the load factor variable γ_2, and its first stage estimates are reported Table A3 in the Appendix. The market competitiveness of neighboring markets has a positive impact on the load factor. To be specific, the load factor of a segment that a product uses increases when a high level of competition exists in its neighboring market (e.g. higher number of nonstop carriers or existence of low cost carrier). This suggests that higher competition in the neighboring market induces more passengers from the market to take the segment's seats, leading to an increase in the load factor of the segment.

[15] I use the following formula (classified in Conlon and Mortimer (2018)) to calculate the diversion ratio of product j to product k in the nested logit setting: $\frac{\partial s_k}{\partial p_j}/|\frac{\partial s_j}{\partial p_j}| = \frac{\sigma s_{k|g}-(1-\sigma)s_k}{1-\sigma s_{j|g}-(1-\sigma)s_j}$, where s_j and $s_{j|g}$ are the market share and the within group market share of product j, respectively, and σ is a nesting parameter.

Table 5: Marginal Cost Estimation Results

	(1) Baseline	(2) Baseline	(3) Frequency in Demand	(4) Frequency in Demand
Load Factor (ν)	0.229***	1.979***	0.289***	2.240***
	(0.001)	(0.006)	(0.001)	(0.0006)
Load Factor (γ_2)	0.975***	0.229***	0.970***	0.204***
	(0.042)	(0.009)	(0.038)	(0.009)
Distance	0.457***	0.246***	0.471***	0.238***
	(0.046)	(0.020)	(0.043)	(0.019)
Distance sq.	0.005	0.022***	0.002	0.022***
	(0.012)	(0.006)	(0.012)	(0.006)
Slot Constraint	−0.085***	−0.039***	−0.090***	−0.036***
	(0.015)	(0.006)	(0.014)	(0.006)
Slot Constraint X LCC	0.282***	0.048**	0.311***	0.035*
	(0.055)	(0.019)	(0.055)	(0.019)
Hub	−0.012	0.015***	−0.014	0.017***
	(0.014)	(0.004)	(0.013)	(0.004)
Base Demand Spec in Table 4?	(1)	(2)	(4)	(5)
O-D Pair (Airport or City)	Airport	Airport	Airport	Airport
Cost Impact of Load Factor (LF)				
LF = 0.6	$ 39.0	$ 6.2	$ 38.5	$ 5.0
(Sample Average) LF = 0.8	$ 43.4	$ 15.8	$ 44.1	$ 14.4
LF = 1.0	$ 45.7	$ 24.6	$ 47.1	$ 23.8
First stage F-Stat (Load Factor)	232883.7	4885.5	151080.9	3884.3
Observations	358,880	358,880	358,880	358,880

Note: $*** < 0.01$, $** < 0.05$, $* < 0.10$. Time, carrier, and non-directional market fixed effects are included in all regressions. The coefficients are reported in units of $100.

6.3 Fixed Cost

Table 6 reports the estimation result for the fixed cost parameters. In the fixed cost specification where only a constant variable is used in column (1), the mean fixed cost ranges from $0.83(M) to $2.91(M). When I add more observable variables to the fixed cost in column (2), the bounds of the parameters are wider, partially because of the Assumption 1 that confidence sets bound of counterfactual segments are wide enough for the purpose of dealing with the selection bias. The constant parameter is positive ranging from $1.76(M) to $3.23(M), but the sets of the rest parameters include zero. Fairly recently, a new working paper that simplifies the confidence set construction for linear conditional moment inequalities (Andrews et al. (2019)) was released. As their approach is similar to my setting, I plan to utilize their approach to compute the bounds of fixed cost parameters which may allow me to estimate more narrow and precise confidence sets.

For the counterfactual analysis, I use $\tilde{\eta}' = (3, -0.7, -1.2, -3)$ in column (2) for three

Table 6: Fixed Cost Parameter Estimation Results

x_{fs}	(1)		(2)	
Constant	[0.837,	2.918]	[1.759,	3.729]
Low Cost Carrier			[−1.976,	1.191]
Slot Ratio at DCA			[−5.370,	1.909]
Presence at Non-DCA			[−8.837,	2.201]

Note: Unit is a million dollar.

reasons: 1) These values are within the range in column (2); 2) none of the moment conditions are violated at the values; and, 3) the values allow the model to fit well with data in terms of entry/exit decisions made by Southwest, JetBlue, and the merging firms.

6.4 Model Fit

In this section, I discuss the second stage model fit (i.e. the pre-merger model fit conditional on flight segment decisions), while the model fit regarding the post-divestitures flight segment entry and exit decisions will be discussed in counterfactual section. For this, I simulate 50 sets of all of the demand and cost variables for "DCA Products" from the estimated distributions. Table 7 shows the model performance, conditional on entry (i.e., second-stage only). The model seems to reasonably predict slot allocation, (weighted) average price, and the number of passengers. However, the model slightly overestimates (underestimates) slots on segments with a lower (higher) slot frequency, and the average price seems to be underestimated for the high-price product group.

To further understand the second stage model fit, I examine the model prediction at an individual flight-segment level. Figure A6 highlights the fact that while the average number of daily slots for a large fraction of flight segments is reasonably well predicted, there is a set of flight segments for which daily slot prediction is not so accurate. Those flight segments with poor predictions, measured by the absolute gap of daily slots between the data and the model being greater than 2 in the figure, mostly contain composite products that are associated with hub airports at which passengers transfer to go to other destinations beyond the perimeter rule. Since the demand qualities of composite products and the sizes of the markets that they belong to are manually chosen to roughly capture the number of connecting passengers, the daily slot gap between data and model predictions for those segments is large. By appropriately assigning demand qualities of the composite products and market sizes of their markets, the second stage model prediction can be further improved.

Table 7: Model Fit (Pre-merger): Slot Allocation, Average Prices and Passengers

	Group	Data	Model Prediction
Slot Allocation			
Daily Departure Frequency Group	1st Tercile	1.53	1.92 (0.11)
	2nd Tercile	3.56	3.53 (0.17)
	3rd Tercile	7.85	7.48 (0.29)
	All	4.31	4.31 (0.19)
Average Price (Weighted by Passengers)			
Data Price Dist. Group	1st Tercile	$156.1	$158.5 (5.5)
	2nd Tercile	$205.4	$196.2 (5.4)
	3rd Tercile	$283.6	$240.7 (5.6)
	All	$196.2	$187.7 (5.5)
Average Passengers			
Data Price Dist. Group.	1st Tercile	5,598.3	5,225.0 (422.7)
	2nd Tercile	3,131.7	3,192.4 (292.8)
	3rd Tercile	2,517.5	2,645.8 (254.1)
	All	3,796.7	3,727.9 (325.8)

I simulate 50 sets of all of the demand and cost variables for "DCA Products" from the estimated distributions. Bootstrapped standard errors are in parenthesis ($n = 40$). For the average prices and passengers, I exclude those composite products.

7 Counterfactuals

In this section, in the context of the AA/US merger, I perform various counterfactual exercises to aid antitrust authorities in measuring the effect of alternative slot divestitures schemes on passenger welfare and market competition *ex ante*. I denote a slot purchaser as a carrier that receives divested slots. It is worth noting that I do not consider the prices that the slot purchaser have to pay for slots in this analysis.

7.1 Counterfactual Description

The merger simulations in this section are based on the demand and marginal cost estimates of "Baseline" specifications in column (2) of Tables 4 and 5 respectively, and on the fixed cost parameters $\tilde{\eta}'$ described in the previous section.

Counterfactual Scenarios

In the context of the AA/US merger case, 104 landing slots that the merging firms had at DCA prior to the merger were divested; these divested slots represented approximately 15% of the average daily flights of the merging firms at the airport.[16] Southwest, JetBlue, and Virgin America ended up purchasing 56, 40, and 8 of the DCA landing slots, respectively, from the merged carrier through a slot auction.[17] The most natural way to measure the effectiveness of the slot divestitures is to compare the simulated merger outcomes in the case in which slot purchasers and the number of divested slots are the same to what actually happened with those in the case in which the merger occurred without slot divestitures. I call this scenario *Baseline*. In this scenario, I also examine to what extent consumer surplus and producer surplus change as the fraction of divested slots increases.

Next, I explore the likely competition effects if slots were granted to different types of carriers (i.e. legacy carriers vs. LCCs). As product demand and the cost structure are heterogeneous across segments and markets, carriers may have different segment entry decisions when additional slots are given to them. This implies distinct welfare effects depending on who receives the divested slots. A particularly interesting case is the AA/US merger, where based on the view that LCCs are more suitable and effective competitors than are legacy carriers, the government approved only LCCs as purchasers of the divested assets. A legacy rival (i.e., Delta) claimed that legacy carriers would bring more services to small and medium-sized communities through its expansive domestic and international network, while the LCCs' business would mainly focus on carrying leisure-based passengers between large domestic destinations—see Gravath, Swaine & Moore LLP (2014). In this counterfactual, I examine merger outcomes when slots are divested to different types of carriers. Additionally, I check if Delta would be likely to enter those flight segments in small communities as it claimed under a scenario in which additional slots are given to Delta. I call this scenario *Purchaser Type*.

Last, I explore the likely competition effect of slot divestitures when the number of slot purchasers varies. Conditional on a fixed number of divested slots, the fraction of slots that each purchaser obtains would be smaller as the number of slot purchasers increases. The likely market structure when the divested slots are evenly split among a few carriers will differ from when they are split among many. However, its prediction is not clear because the flight segment entry decisions differ by carriers. In this exercise, I compare

[16] In this merger deal, the merged carrier was also required to divest 34 landing slots at LaGuardia (LGA) and some airport gate access at Reagan National (DCA), LaGuardia (LGA), Boston Logan (BOS), Chicago O'Hare (ORD), Dallas Love Field (DAL), and Los Angeles (LAX).

[17] Of the 40 acquired landing slots, JetBlue won 24 via the auction, and the remaining 16 were obtained by permanently controlling the slots that had been leased to JetBlue before the merger by American Airlines.

Table 8: List of Marginally Profitable Flight Segments from DCA

	Entry				Exit
Rank	Southwest	JetBlue	Delta	United	NewAA
1	FLL	PBI	MIA	PNS	CAK
2	TPA	SRQ	BOS	FLL	AGS
3	BNA	BDL	FLL	TYS	MYR
4	MSY	JAX	MCO	PVD	TLH
5	MCI	RSW	JAX	SAV	JAN
6	PVD	CHS	IND	MIA	PNS
7	CMH	MSY	BDL	DSM	GSP
8	BHM	PWM	PVD	TPA	BNA
9	JAX	CLT	OMA	DAY	ROC
10	IND	MSP	MSY	JAX	TYS

A simple logit model of the actual entry decision made by Southwest and JetBlue after the AA/US merger is used to predict the likelihood of entry/exit. The logit estimate can be found in Appendix C.1. Additionally, the three-letter airport codes used in this table are defined by the International Air Transport Association (IATA). Table A4 shows the full airport names of those codes.

merger outcomes by gradually increasing the number of purchasers from one to four. I call this scenario *Number of Purchasers*.

Marginally Profitable Flight Segments

While the model allows carriers to endogenously choose flight segments, it is computationally infeasible to compute the expected profits for all possible combinations of flight segments. For example, if there are 60 different flight segments from DCA, 2^{60} sets of endogenous routes are possibly chosen by each carrier, and it is computationally very expensive to compute equilibrium slot allocations and prices for all of the combinations. To overcome this, I assume that carriers can endogenously choose only a subset of flight segments that are marginally profitable. The merged carrier that has a reduced number of landing slots following the slot divestitures has an incentive to remove from its airline network the existing flight segments that make a relatively small contribution to the carrier's profit. Analogously, carriers that purchase additional landing slots from the divestitures could add new flight segments to their networks that were previously considered not as profitable because those segments already exist in their networks.

To obtain marginally profitable flight segments, based on a simple logit model, I predict the likelihood of new entry by using the information of the post-merger entry decisions made by the actual slot purchasers in the AA/US merger—Southwest and JetBlue.[18] The expected variable profit change when a carrier adds a counterfactual segment

[18]While I use the *ex post* information in this *ex ante* analysis, we can build a similar logit model by

is included as a main explanatory variable in the model. Additionally, a carrier's airport presence at a counterfactual segment's endpoints is included as the variable plays an important role in entry decisions—see Goolsbee and Syverson (2008). While the detailed prediction procedure is available in Appendix C.1, Table 8 lists by carrier the top ten flight segments that are most likely to be added/removed from DCA. The three-letter airport codes defined by the International Air Transport Association (IATA) are used in the Table, and their full names can be found in Table A4 in the Appendix.

I make several post-merger assumptions. First, I compute the counterfactual outcomes based on the assumption that NewAA takes the average observed characteristics of the two merging firms pre-merger. As Ciliberto et al. (2018) point out, I recognize that merger simulation outcomes may change substantially, depending on our assumptions about the characteristics of the merging firms post-merger. In this counterfactual exercise, however, as we are more interested in the relative differences between various divestiture schemes than in the absolute surplus measures for each scenario, I make a single assumption about the quality of NewAA. Second, the entry game in this counterfactual analysis assumes that the order of moves is an equilibrium selection mechanism. In the endogenous entry game setting, there are concerns over multiple equilibria when carriers simultaneously make entry decisions. Following Wollmann (2018) and Lee and Pakes (2009), to alleviate this concern, I sort the carriers based on their airport passenger shares at DCA and assume that the carrier with the highest passenger share moves first in the sequential game, followed by the second highest one and so forth. Last, I assume that the segment-specific airplane size post-merger is the same as the one pre-merger. There might be a concern that NewAA could change its fleet allocation after the merger, which could systematically alter the segment-specific airplane size. However, a regression exercise in Appendix C.2 suggests that this did not occur.

Equilibrium: Brute Force Search vs. Heuristic Search

I take two approaches to find the set of flight segments that carriers choose at the equilibrium. One approach, which I call a "Brute Force Search," is to compute the expected profit for every possible combination of endogenous flight segments. While the "Brute Force Search" guarantees the accuracy of the model solution within the choice of the set of endogenous flight segments, it becomes computationally intensive as the number of endogenous segments increases. For example, if carriers were allowed to select all the segments in the list of Table 8, there would be 2^{50} combinations to analyze (each carrier has 2^{10} combinations), which is an examination that is computationally infeasible. For this search, therefore, I further refine a set of marginally profitable routes from the list in

leveraging the information of the previous events such as historical mergers with slot divestitures or slot swap events.

Table 8. I describe the refinement in the following subsections.

The other approach, a "Heuristic Search," allows the equilibrium to be found quickly with a relatively large number of endogenous segments. Similar to the approach Fan and Yang (2016) take in their counterfactual analysis, this approach is based on the iterated best response in which each carrier takes a turn and grows/shrinks a set of endogenous segments until they reach to the point where there is no profitable deviation for all carriers. In contrast to the "Brute Force Search," carriers are allowed to grow/shrink the set of segments by up to one flight segment for each turn in this search. This potentially implies that flight segment composition can vary if carriers are allowed to add/remove more than one flight segments. In Appendix C.3, I describe how the "Heuristic Search" works in detail.

7.2 Baseline

In this scenario, I compare the simulated merger outcomes in the case in which slot purchasers and the number of divested slots are the same as in the case in which the merger occurred without slot divestitures. I assume that Southwest and JetBlue take 60% and 40% of divested slots, respectively, based on the ratio of slots that they actually obtained through the divestiture process. I exclude Virgin America in this analysis as it did not operate any nonstop flights within the perimeter rule at DCA pre-merger. For the "Brute Force Search," I choose the first four, five, and four endogenous segments to/from DCA for NewAA, Southwest, and JetBlue, respectively, in Table 8. There were two overlapping markets from/to DCA where the merging firms were duopolists—Nashville (DCA⇒BNA) and Raleigh (DCA⇒RDU). To assess whether slot purchasers had an incentive to constrain the market power of the merged firm in those overlapping markets, I allow Southwest to endogenously choose the two flight segments.[19] Figure 4 visualizes the list of endogenous segments marked as a dotted black line, while the existing exogenous segments are marked as a solid gray line.

Table 9 shows the result of the simulated merger outcomes of *Baseline* case under the "Brute Force Search" approach. The simulation results show that the average price would increase and that the number of passengers would decrease when there were no slot divestitures. As the number of divested slots increase, DCA markets become more competitive and consumers would be better off. The effect of slot divestitures on the surplus is substantial. For example, the consumer surplus gap between the 15% slot divestitures case and no divestiture is computed as \$28.1(M) per quarter or roughly \$112.4(M) per year.

[19] As Southwest obtained more than 60% of the divested slots when excluding those slots leased to JetBlue, the carrier had greater flexibility to enter new routes than other LCCs.

Figure 4: Marginally Profitable Routes by Carrier (*Baseline*)

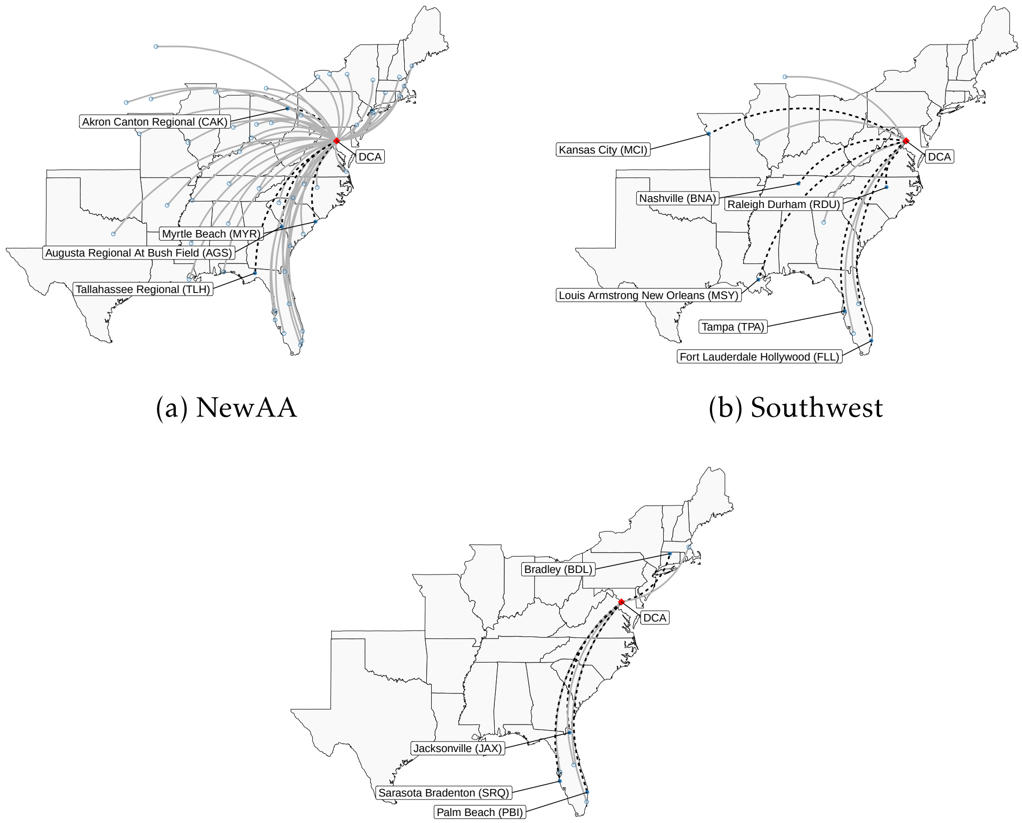

(a) NewAA

(b) Southwest

(c) JetBlue

Note: Gray solid lines indicate the flight segments treated as exogenous, and black dotted lines indicate those segments that carriers endogenously choose. Those endogenous segments are labeled on the maps.

The merger simulation suggests that the merger remedy will have distributional effects across markets. To understand this, I decompose the consumer surplus into three categories in the second panel of the table—1) "No Change" refers to the markets where there are no product entries or exits after the merger; 2) "Segments Added" is a group of markets in which new products are introduced due to new segment entries; and 3) "Segment Removed" is a group of markets in which the existing products are removed as carriers exit the corresponding flight segments. As the ratio of divested slots increases, slot purchasers initiate new nonstop services in new markets and passengers in those markets are better off due to intense market competition. This is explained by the fact that the consumer surplus change in "Segments Added" is greater than the one in "No Change."

However, slot divestitures are not necessarily good for everyone. For example, consumers in the "Segment Removed" markets will be negatively affected by the divestitures. The products of NewAA in those markets are no longer available as the firm exits the mar-

Table 9: Post-merger Outcomes (*Baseline*) Using Brute Force Search

	Pre-merger	Post-merger (Slot Divestitures Ratio)			
		0%	10%	15%	20%
All DCA Markets					
Price	$180.26	+5.25%	-0.37%	-1.66%	-2.78%
Passengers	2,856(k)	-2.91%	+0.42%	+0.91%	+1.45%
Consumer Surplus		-25.51(M$)	+0.07(M$)	+2.60(M$)	+5.35(M$)
All DCA Markets Consumer Surplus Decomposition (M$)					
No Change		-28.76	-18.11	-17.80	-13.74
Segments Added		+4.03	+22.09	+25.12	+24.32
Segments Removed		-0.77	-1.32	-2.13	-2.66
Overlapped Markets					
Price	$249.89	+37.47%	+20.34%	+19.28%	-9.76%
Passengers	62(k)	-36.63%	-25.77%	-25.45%	+2.56%
Consumer Surplus		-6.71(M$)	-4.72(M$)	-4.66(M$)	+0.47(M$)

Note that units in Post-merger columns are all relative to Pre-merger values. Overlapped markets to/from DCA include DCA⇒BNA, BNA⇒DCA, DCA⇒RDU, and RDU⇒DCA markets. I use 40 draws to obtain the expected profit for each combination of flight segments.

kets and redistributes its relatively scarcer slots to other more profitable segments. This leads to a consumer surplus loss in those markets. Additionally, there is no guarantee that slot purchasers will enter segments to serve the overlapping markets(i.e., the markets in which the merging firms were duopolists pre-merger). In the last panel of the Table 9, the consumer surplus in the overlapped markets tends to stay negative. This is because Southwest would prioritize entering more profitable segments over entering the segments related to the overlapped markets. The model predicts that the carrier would enter all the overlapped markets when it has abundant slots (20% slot divestiture case).

The potential concern that slot purchasers may not serve overlapped markets motivates an alternative merger policy. What if, as a condition for buying slots, the antitrust authorities require a slot purchaser to serve a set of flight segments? Table 10 shows the comparison of consumer/producer surplus in the case in which Southwest is required to enter those overlapped markets vs. the case in which it is not required to enter the markets. Compared to the Baseline (B) where there is no government intervention (same numbers in the *Baseline* case), when Southwest is forced to enter the Nashville and Durham segments, there are consumer surplus gains and producer surplus losses. Consumers in the overlapped markets will be better off in this intervention. In terms of producer surpluses, as the slot purchaser would have been more profitable if it were freely choosing segments, its profit under the intervention would be lower than in the baseline case.

Table 10: Surplus Changes When a Purchaser is Forced to Enter Overlapped Markets

	Consumer Surplus (M$)			Producer Surplus (M$)		
divest	Baseline (B)	Forced (F)	(F)-(B)	Baseline (B)	Forced (F)	(F)-(B)
All DCA Markets						
10%	0.07	3.78	3.71	33.19	31.09	-2.10
15%	2.60	6.68	4.08	30.89	28.87	-2.02
20%	5.35	5.35	0.00	27.86	27.86	0.00
Overlapped Markets						
10%	-4.72	0.10	4.82			
15%	-4.66	0.32	4.99			
20%	0.47	0.47	0.00			

Note: 'Baseline (B)' shows merger outcomes without government intervention (*Baseline*), while 'Forced (F)' shows merger outcomes when Southwest is forced to enter the overlapped markets (DCA⇔RDU and DCA⇔BNA). The values in those columns indicate the post-merger surplus relative to the pre-merger one. I use 40 draws to obtain the expected profit for each combination of flight segments.

Next, I examine the extent to which surpluses change as the amount of divested slots increases in Figure 5 using the heuristic approach. For this figure, ten endogenous flight segments for NewAA, Southwest and JetBlue shown in Table 8 are used. For each set of 20 draws of demand and marginal cost unobservables, I calculate the expected profit, and 25 sets are used to obtain the confidence intervals of surpluses, as each set may have a different equilibrium flight segment choice.

As the slot divestiture ratio increases, consumer and produce surpluses tend to go in the opposite directions. On the one hand, consumer surplus tend to increase in slot divestiture as in Figure 5 because of those new flight segments initiated by slot purchasers and of the existing products getting cheaper from a relaxed capacity constraint. The increase in market competition allows consumers to be better off. On the other hand, the model predicts that the NewAA would have less seats and its capacity constraint is likely to be binding, as there are more slots to be divested. This leads to a profit loss of the NewAA, and this loss outweighs the profit gains by slot purchasers. Last, the total surplus, i.e., the sum of the consumer and producer surpluses, is provided in this analysis which helps understand the merger effect as a social planner's point of view. The total surplus in Figure 5, marked as the blue square, increase but is saturated for the higher ratio of slot divestitures.

My model fits the data well in terms of the post-divestiture flight segment entry decisions. In Table A5, I list the actual new flight segments offered by Southwest and JetBlue post-divestitures (in regular font face) and the model predicted new segments (in bold font face) when using the heuristic search under the 15% slot divestitures scenario (the realized divestiture scenario). The table shows that, first, the list of new segments in the data highly overlaps with the list of marginally profitable flight segments in Table 8 (e.g.

Figure 5: Post-merger Outcomes (*Baseline*) Using Heuristic Search

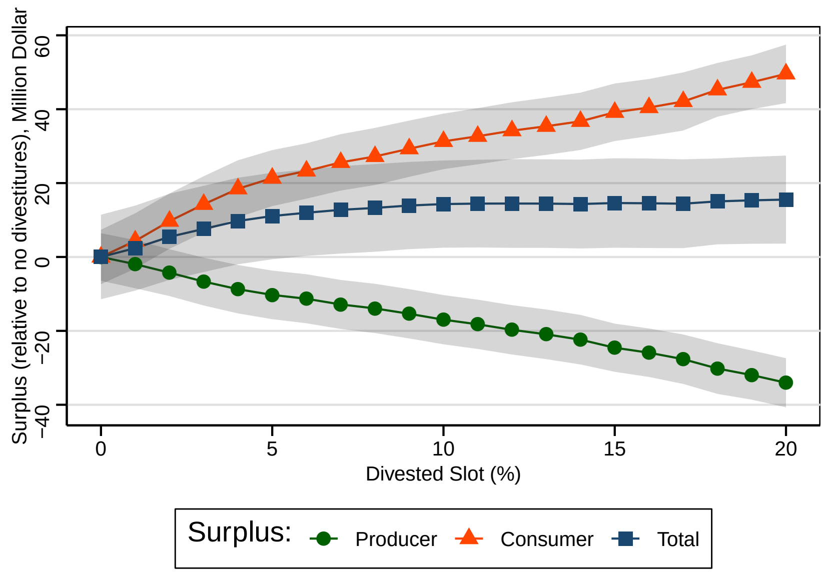

I use 20 draws for each set to compute the expected variable profits and to find the equilibrium. Then, I use 25 sets to calculate the 95% confidence interval of surpluses.

the matching rate is 80% for Southwest and 100% for JetBlue, respectively). Second, while there are more new segments in the data than what the model predicts, the proportion of correctly matching the new entry segments is high. For example, the proportion of the number of the model-predicted segments to the number of actual segments for Southwest and JetBlue is 70% and 60%, respectively. In terms of the flight segment exit decisions, we see clearly in the data that NewAA eliminated a set of small community-based flight segments post-divestiture (e.g. AGS, LIT, MYR, OMA, and TLH). However, the LIT and OMA segments are not considered endogenous segments in Table 8, and MYR is the only segment that the model correctly predicted. Potentially, the existence of commuter slots and how NewAA allocates those slots to small community-based segments can be attributed to this discrepancy between the model and data, given that the actually divested slots are regular not commuter slots.

7.3 Purchaser Type

While slots were solely granted to LCCs in *Baseline*, I explore the case where legacy carriers were considered as slot purchasers as a comparison. To do so, I assume that Delta and United take 60% and 40% of the divested slots and allow them to endogenously choose up to the first six and four flight segments in Table 8, respectively. Under the "Brute Force Search," Table 11 compares the merger simulation outcomes by slot purchaser types. When additional slots were granted to them, carriers have heterogeneous

Table 11: Post-merger Outcomes (*Purchaser Type*) Using Brute Force Search

Purchaser Type?	LCCs		Legacy Carriers	
Slot Divestitures Ratio?	10%	15%	10%	15%
All DCA Markets				
Price	-0.37%	-1.66%	0.66%	-0.04%
Passengers	0.42%	0.91%	-0.57%	-0.50%
Consumer Surplus	0.07(M$)	2.60(M$)	-6.97(M$)	-7.89(M$)
List of Flight Segments Added(+)/Removed(-)				
Merging Firm (-)	CAK MYR	CAK MYR TLH	CAK MYR	CAK MYR TLH
Southwest (+)	FLL TPA BNA MSY	FLL TPA BNA MSY MCI		
JetBlue (+)	PBI BDL JAX	PBI BDL JAX		
Delta (+)			MIA BOS FLL	MIA BOS FLL
United (+)				

Note that units in Post-merger columns are all relative to Pre-merger values. Overlapped markets to/from DCA include DCA⇒BNA, BNA⇒DCA, DCA⇒RDU, and RDU⇒DCA markets. The full airport names of those three-letter airport codes used in this table can be found in Table A4.

preference on adding new flight segments, depending on their segment level demand and cost characteristics.

The model result does not suggest that the legacy carriers would serve small- or medium-sized communities from/to DCA, in contrast to Delta's claim. This is largely because the model predicts that the carrier would earn more profit by serving popular destinations such as Boston and Miami rather than small-sized communities. Interestingly, United would increase the frequency of the existing segments rather than open new segments, as its fixed cost to open a new segment is high. In addition, the results of the table suggest that for slot divestitures of both 10% and 15%, the overall consumer surplus under LCC purchasers are calculated to be larger than that under legacy purchasers.

The tendency of a legacy carrier to choose those segments related to big cities or leisure-based destinations over those segments related to small- medium-sized communities can be seen more clearly when using the heuristic approach as in Figure 6. For this figure, I assume that the NewAA's divested slots (15%) go to Delta, allowing the carrier to endogenously choose five big community-based segments from DCA (MCO, ORD, BOS, MIA, and FLL) and five small/medium community-based segments (IND, CMH, MKE, JAX, and BDL) that the carrier said it would be willing to enter if it were a slot purchaser in its complaint Gravath, Swaine & Moore LLP (2014). The left panel of the figure shows the surplus change as the amount of divested slots increases. The total surplus relative

Figure 6: Expected Post-merger Outcomes (*Purchaser Type*) Using Heuristic Search

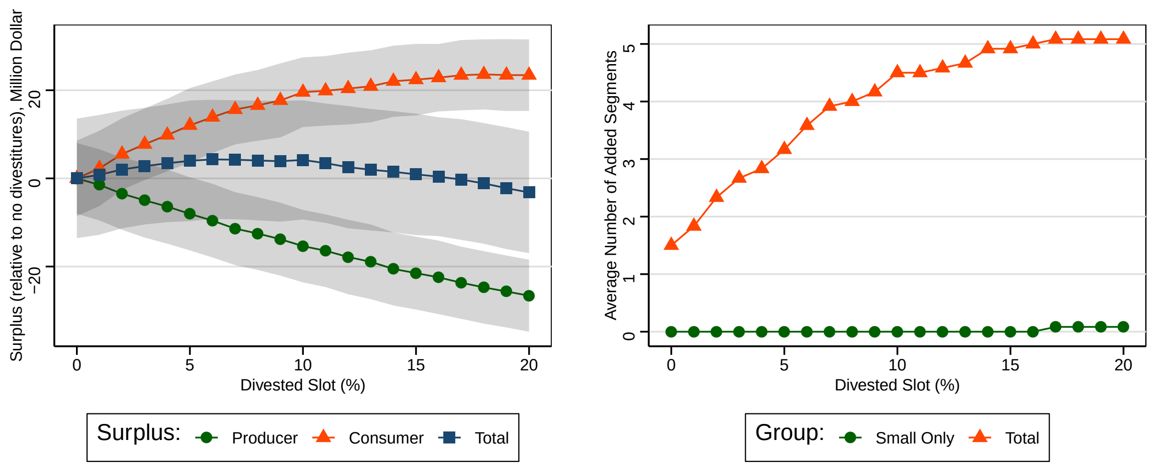

(a) Surplus Changes

(b) Average Number of Added Segments

Note: I use 20 draws for each set to compute the expected variable profits and to find the equilibrium. The 95% confidence interval of surpluses in (a) is based on 25 sets of draws.

to the no divestitures case increases first but eventually decreases and goes below zero when the ratio of divested slots is high. In the right panel, the average number of newly added segments by Delta is reported. While the number of new segments increases in the ratio of slot divestitures, those new segments are mostly big community-based segments and the line "Small Only"—the average number of newly added small/medium community-based segments — is near zero for each divestiture considered.

7.4 Number of Purchasers

In this alternative divestiture scheme, I vary the number of slot purchasers, keeping the divested slots at 15% of the merging firms' total endowed slots, and assume that carriers as slot purchasers have the following orders of modifying their flight segment portfolio—Southwest, JetBlue, Delta, and United. While the number of endogenous flight segments of NewAA is kept as four (e.g. CAK, MYR, AGS, and TLH), the sets of endogenous segments of carriers other than NewAA vary in the number of slot purchasers—1) in the one-slot purchaser scenario, Southwest can choose ten segments in Table 8; 2) in the two-slot purchaser scenario, Southwest can choose five segments, and JetBlue can choose five segments; 3) in the three-slot purchaser scenario, Southwest can choose four, JetBlue three, and Delta three; and, 4) in the four-slot purchaser scenario, Southwest can choose three, JetBlue two, Delta two, and United two. Additionally, I assume that divested slots are evenly split among purchasers. For example, if there are three slot purchasers (Southwest, JetBlue, and Delta), each carrier takes a third of the divested slots.

Table 12: Post-merger Outcomes (*Number of Purchasers*) Using Brute Force Search

Number of Purchasers?	One	Two	Three	Four
All DCA Markets				
Price	-2.21%	-1.66%	-2.31%	-2.20%
Passengers	1.19%	0.91%	1.34%	1.07%
Consumer Surplus	6.20(M$)	2.60(M$)	7.40(M$)	5.02(M$)
List of Flight Segments Added(+)/Removed(-)				
Merging Firm (-)	CAK MYR	CAK MYR	CAK MYR	CAK AGS
	TLH	TLH	TLH	TLH
Southwest (+)	FLL TPA	FLL TPA	FLL TPA	FLL TPA
	BNA MSY	BNA MSY	MSY	
	RDU PVD	MCI		
	CMH JAX			
JetBlue (+)		PBI BDL JAX	PBI BDL	PBI BDL
Delta (+)			MIA BOS	MIA BOS
United (+)				

Note that the units in Post-merger columns are all relative to the Pre-merger values. The number of endogenous segments for each column is the following. i) 'One': Southwest 10 segments; ii) "Two": Southwest 5, and JetBlue 5; iii) "Three": Southwest 4, JetBlue 3, and Delta 3; and iv) "Four": Southwest 3, JetBlue 3, Delta 2, and United 2.

Table 12 shows the simulated merger outcomes by varying the number of purchasers. In this result, the likely consumer surpluses (relative to pre-merger) in the four cases differ with a range of \$2.6(M) to \$7.4(M). More importantly, varying the number of slot purchasers leads to a distributional welfare effect on passengers, as slot purchasers choose different sets of endogenous flight segments. If a few carriers buy the divested slots, there will be a reduction in fixed cost due to an increase in their endowed slot ratio at DCA, facilitating their entry into new flight segments. On the other hand, when more carriers become slot purchasers, carriers will focus on fewer segments where carriers find it the most profitable.

8 Conclusion

I have developed an endogenous entry model that features the carriers' flight segment entry decisions, product price choices and slot allocation choices at a slot-controlled airport. In this model framework, the marginal costs increase as a carrier's load factor on a flight segment increases, enabling a change in capacity of a segment to affect product costs in a large number of markets. Then, focusing on the AA/US merger in which the parties were required to divest a significant proportion of their landing slots at DCA, my paper assesses various alternative slot divestiture schemes at DCA, involving both different numbers of slots and different allocations of the divested slots to rival carriers. I find that the proposed divestiture raised consumer surplus significantly ($112M per year) relative to approving the merger without divestiture, but that it harmed a subset of DCA passengers as the merged firm dropped services on some marginally profitable routes. I also find that the policy of only allowing the slots to be divested to low-cost carriers raised consumer surplus relative to the policy only allowing the slots to be divested to legacy carriers, and that, contrary to some claims made at the time, divestitures to legacy rivals would not have protected services to smaller airports.

The model I developed can potentially be extended not only to non-merger issues in the airline industry, but also to merger remedies in other industries. One example is slot swap—exchanges in which airlines that are dominant at slot-controlled airports trade slots with one another in an attempt to increase their competitive advantage (e.g. DL/US in 2011). Antitrust authorities often find these transactions to be anticompetitive, as slot swapping may further enable the carriers involved to raise prices at their dominant airports. Simply by extending the scope of the airline network from one slot-controlled airport to two, my model can be applied to analyze this policy. In terms of the merger remedy analysis in other industries, retail mergers are an example where divestitures of stores or distribution centers are commonly used. As those assets affect product costs in multiple local markets via transportation costs or inventory controls, my model can be naturally extended to those merger remedies.

Some limitations of my model provide an agenda for future research. While I believe the key features in my model—capacity choice via airport slots and cross-market interactions—extend the existing frameworks for merger simulation and airline models, the model faces a few limitations. First, the consumer demand in the model is not directly affected by divested assets due to model tractability. However, if the model can account for the relationship between divested assets and market demand (e.g. daily departure options in airline mergers or shopping distance to stores/branches in retail or banking mergers), the analysis will be richer. Second, my model does not incorporate the carrier's fleet assignment decision. While this model assumes that a segment-specific airplane size is exogenously given, carriers may change their post-merger fleet portfolio in the long run. Last, I examined only a subset of products and markets in the U.S. airline industry that is associated with the slot divestiture in the AA/US merger. Although the model feature in which markets are interconnected with each other is a clear extension of the existing literature, the scope of airline networks in this study is limited to a slot-controlled airport.

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.

Andrews, D. and Soares, G. (2010). Inference for parameters defined by moment inequalities using generalized moment selection. *Econometrica: Journal of the Econometric Society*, 78(1):119–157.

Andrews, I., Roth, J., and Pakes, A. (2019). Inference for Linear Conditional Moment Inequalities. *Working Paper*.

Ashenfelter, O., Hosken, D., and Weinberg, M. (2014). Did Robert Bork Understate the Competitive Impact of Mergers? Evidence from Consummated Mergers. *The Journal of Law and Economics*, 57(S3):S67–S100.

Ball, M. P., Ausubel, L. M., Berardino, F., Cramton, P., Donohue, G., Hansen, M., and Hoffman, K. (2007). Market-Based Alternatives for Managing Congestion at New York's LaGuardia Airport. pages 1–31.

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*, 63(4):841.

Berry, S. T. (1992). Estimation of a Model of Entry in the Airline Industry. *Econometrica: Journal of the Econometric Society*, 60(4):889–917.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND journal of economics*, 25(2):242.

Besanko, D. and Doraszelski, U. (2004). Capacity Dynamics and Endogenous Asymmetries in Firm Size. *The RAND journal of economics*, 35(1):23.

Borenstein, S. (1990). Airline mergers, airport dominance, and market power. *American Economic Review*, 80:400–404.

Bresnahan, T. F. and Reiss, P. C. (1990). Entry in Monopoly Markets. *The Review of Economic Studies*, 57(4):531–553.

Ciliberto, F., Murry, C., and Tamer, E. (2018). Market Structure and Competition in Airline Markets. *Working Paper*.

Ciliberto, F. and Tamer, E. (2009). Market Structure and Multiple Equilibria in Airline Markets. *Econometrica: Journal of the Econometric Society*, 77(6):1791–1828.

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

Conlon, C. T. and Mortimer, J. H. (2018). Empirical Properties of Diversion Ratios. *National Bureau of Economic Research Working Paper Series*.

Duso, T., Gugler, K., and Yurtoglu, B. (2007). EU merger remedies: An empirical assessment. *The Political Economy of Antitrust*, 282(4):303–348.

Eizenberg, A. (2014). Upstream Innovation and Product Variety in the U.S. Home PC Market. *Review of Economic Studies*, 81(3):1003–1045.

Fan, Y. (2013). Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market. *American Economic Review*, 103(5):1598–1628.

Fan, Y. and Yang, C. (2016). Competition, Product Proliferation and Welfare: A Study of the U.S. Smartphone Market. *SSRN Electronic Journal*.

FTC (1999). A study of the Commission's divestiture process. Technical report.

FTC (2017). The FTC's Merger Remedies 2006-2012. Technical report.

Goolsbee, A. and Syverson, C. (2008). How Do Incumbents Respond to the Threat of Entry? Evidence from the Major Airlines. *The Quarterly Journal of Economics*, 123(4):1611–1633.

Gravath, Swaine & Moore LLP (2014). Comments Of Delta Air Llines, INC. Concering Proposed Final Judgement As To defendants US Airways Group, INC. and AMR CORP. . pages 1–37.

Hatzitaskos, K., Kirgiz, K., michelozzi, L., and Hoffman, I. (2019). Trends in Merger Investigations and Enforcement at the U.S. Antitrust Agencies. Technical report.

Ho, K. and Rosen, A. (2015). Partial Identification in Applied Research: Benefits and Challenges. Technical report, National Bureau of Economic Research, Cambridge, MA, Cambridge, MA.

Holmes, T. J. (2011). The Diffusion of Wal-Mart and Economies of Density. *Econometrica: Journal of the Econometric Society*, 79(1):253–302.

Hüschelrath, K. and Müller, K. (2014). Market Power, Efficiencies, and Entry Evidence from an Airline Merger. *Managerial and Decision Economics*, 36(4):239–255.

Israel, M., Keating, B., Rubinfeld, D. L., and Willig, B. (2013). Airline Network Effects and Consumer Welfare. *Review of Network Economics*, 12(3):1–36.

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

Kwoka, J. (2014). *Mergers, Merger Control, and Remedies*. A Retrospective Analysis of U.S. Policy. MIT Press.

Kwoka, J. (2017). Merger Remedies: An Incentives/Constraints Framework. *The Antitrust Bulletin*, 62(2):367–381.

Kwoka, J. and Shumilkina, E. (2010). The Price Effect Of Eliminating Potential Competition: Evidence From An Airline Merger. *The Journal of Industrial Economics*, 58(4):767–793.

Lee, R. S. and Pakes, A. (2009). Multiple equilibria and selection by learning in an applied setting. *Economics Letters*, 104(1):13–16.

Li, S. Y., Mazur, J., Park, Y., Roberts, J., Sweeting, A., and Zhang, J. (2019). Repositioning and Market Power After Airline Mergers. *Working Paper*, pages 1–83.

Nevo, A. (2001). Measuring Market Power in the Ready-to-Eat Cereal Industry. *Econometrica: Journal of the Econometric Society*, 69(2):307–342.

Pakes, A., Porter, J., Ho, K., and Ishii, J. (2015). Moment Inequalities and Their Application. *Econometrica*, 83(1):315–334.

Reiss, P. C. and Spiller, P. T. (1989). Competition and Entry in Small Airline Markets. *The Journal of Law and Economics*, 32(2).

Reitzes, J. D., McVeigh, B., Powers, N., and Moy, S. (2014). Competitive Effects of Exchanges or Sales of Airport Landing Slots. *Review of Industrial Organization*, 46(2):95–125.

Seim, K. (2006). An empirical model of firm entry with endogenous product-type choices. *The RAND journal of economics*, 37(3):619–640.

Snider, C. (2009). Predatory Incentives and Predation Policy: The American Airlines Case. pages 1–48.

United States (2013). Complaint : U.S. and Plaintiff States v. US Airways Group, Inc. and AMR Corporation. *Department of Justice*, pages 1–56.

United States (2019). Oversight of the Enforcement of the Antitrust Laws. *Department of Justice*, pages 1–21.

Villas-Boas, S. B. (2007). Vertical Relationships between Manufacturers and Retailers: Inference with Limited Data. *The Review of Economic Studies*, 74(2):625–652.

Wang, W. and Rudanko, M. (2012). EU Merger Remedies and Competition Concerns: An Empirical Assessment. *European Law Journal*, 18(4):555–576.

Wollmann, T. G. (2018). Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles. *American Economic Review*, 108(6):1364–1406.

Appendix

A Figures and Tables

Figure A1: Hourly Slot Holdings By Carrier Type Before/After Slot Divestitures At DCA

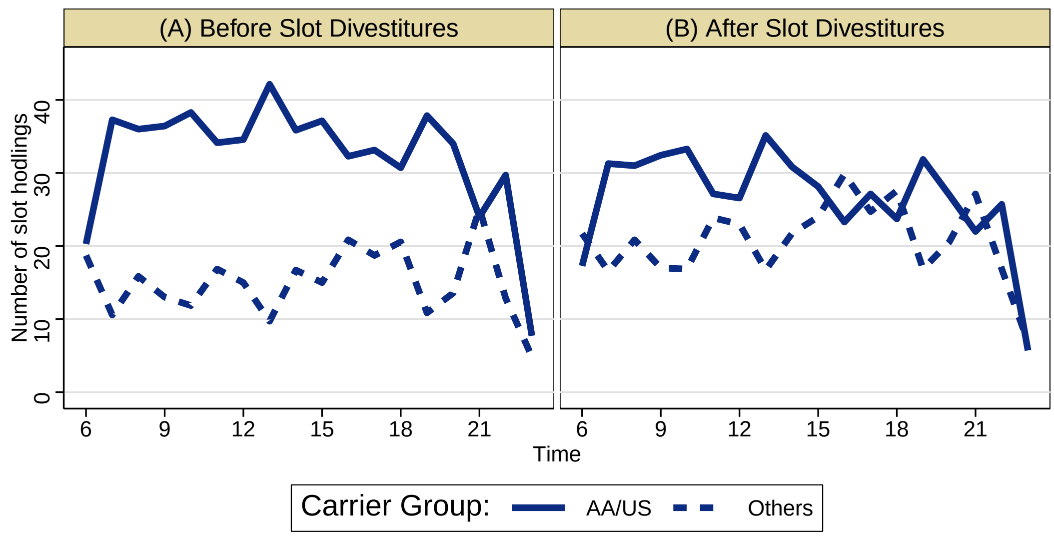

Note: The figures show the number of slot holdings of merging firms (marked as solid line) and other carriers (marked as dotted line) before and after the slot divestitures at DCA. The information on slot holdings is extracted from the Slot Administration page of the Federal Aviation Administration website.

Figure A2: Histogram of Average Daily Slots Assigned on Segments at DCA (2013Q2)

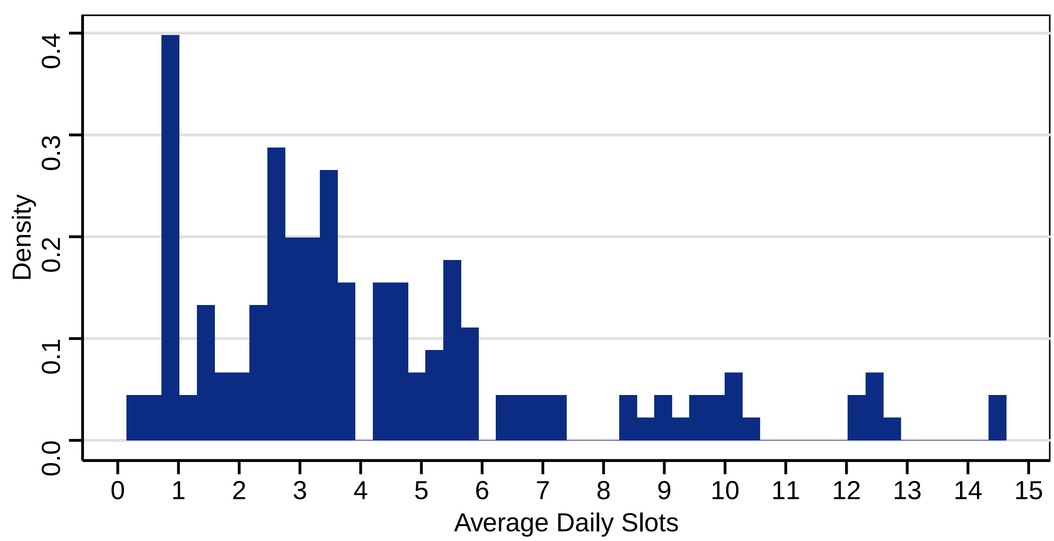

Note: This histogram is based on the average daily slots of DCA flight segments in 2013Q2 (n=156).

Figure A3: Number of Slot Holdings and Average Number of Flight Operations at DCA (2018Q2)

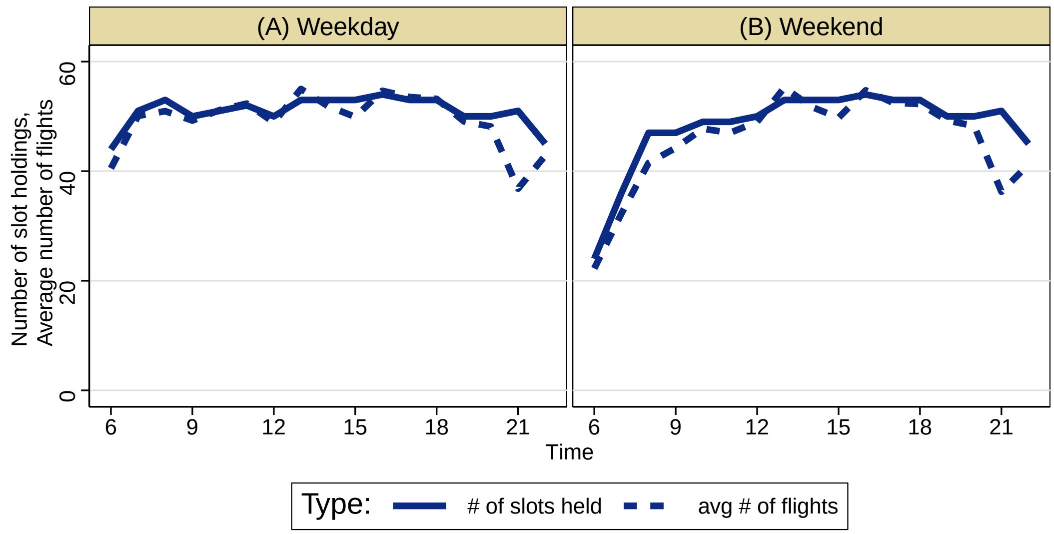

Note: The number of slots held by any commercial carrier in each time bin (one hour block) as of June 2018 is displayed in a solid line, and the average number of scheduled flights in each time bin (operated by any carriers) is displayed in a dashed line. The information on slot holdings is extracted from the Slot Administration page of the Federal Aviation Administration website, and the information on the average number of flight operations is based on the Marketing Carrier On-Time Performance Data from Bureau of Transportation Statistics.

Figure A4: Illustration of Composite Connecting Products

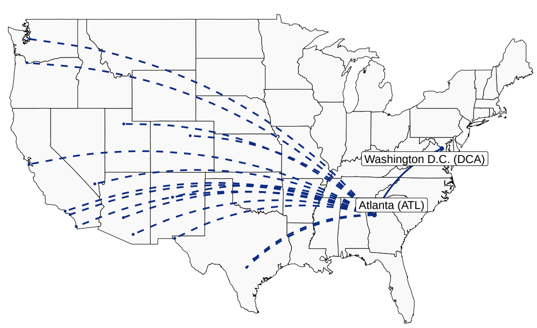

(a) DL's Composite Product

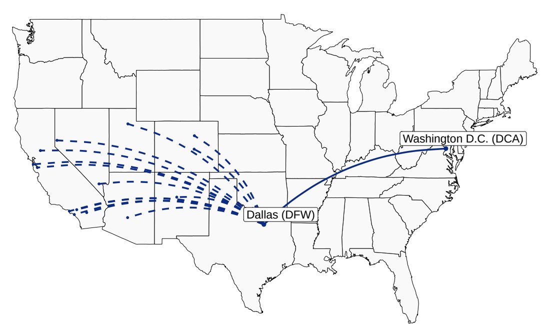

(b) AA's Composite Product

Note: The first panel shows 14 different itineraries (products) offered by Delta when connecting at Atlanta (ATL) to go to DCA. The nonstop distances of the corresponding markets of those itineraries are beyond the perimeter rule threshold (1,250 miles), while $(\text{DCA}\rightarrow\text{ATL})_{DL}$ is within the perimeter rule. I combine those 14 products into a single composite product in order to take into account their small but non-negligible effect on the load factor of $(\text{DCA}\rightarrow\text{ATL})_{DL}$. Analogously, the second panel shows a set of connecting products offered by American connecting at Dallas (DFW). To generate a moderate level of market competition, I assume that the two composite products are in the same market.

Figure A5: Empirical Distribution of Demand and Marginal Cost Residuals

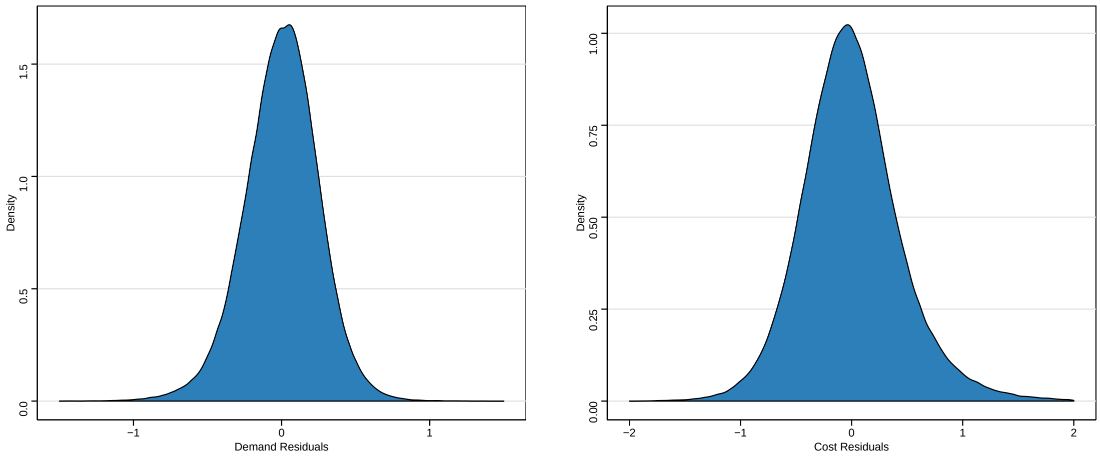

(a) Demand Residuals ($\hat{\xi}$) (b) Marginal Cost Residuals ($\hat{\omega}$)

Note: The empirical distributions of demand residuals (left) and marginal cost residuals (right) are reported (Airport-Airport pair baseline assumption). One standard deviation of demand and marginal cost residuals are equal to \$74.23 and \$25.35, respectively. In addition, note that the average marginal cost is computed as \$73.46.

Figure A6: Average Number of Daily Slots (Data and Model Prediction)

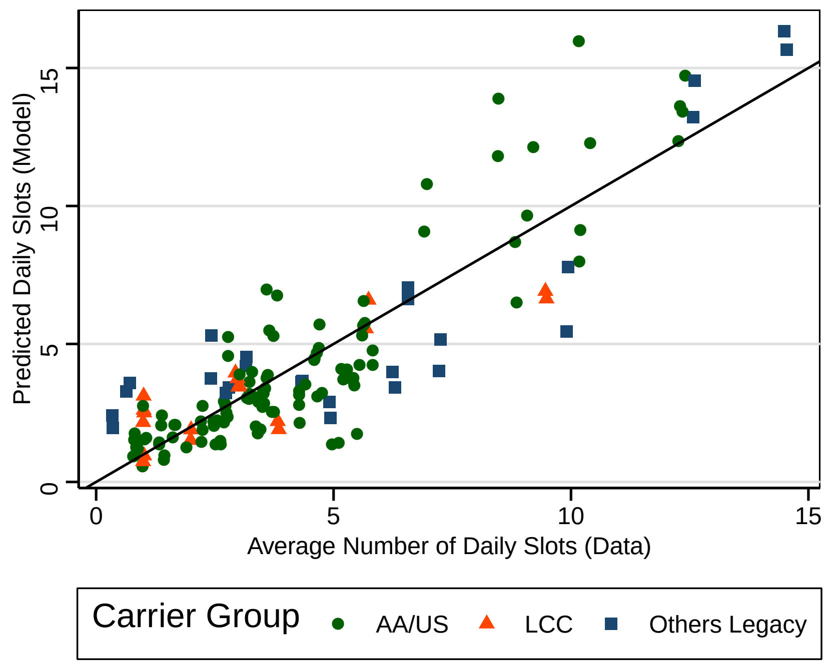

Note: This scatter plot compares the average number of daily slots from the data (on the x-axis) with the average number of daily slots from the model prediction (on the y-axis). A 45 degree line is marked as a solid black line.

Table A1: Regression Results of Airplane Size

	Dependent variable:		
	Airplane Size		
	(1)	(2)	(3)
Distance (1,000 miles)	43.326***	41.659***	20.157***
	(4.055)	(4.145)	(3.183)
Distance^2	−3.385**	−2.631*	0.041
	(1.517)	(1.573)	(1.197)
International Hub	21.840***	22.622***	
	(2.240)	(2.248)	
Slot-Controlled (All)		−8.460***	
		(2.002)	
Origin-Carrier F.E.?	No	No	Yes
Dest-Carrier F.E.?	No	No	Yes
Observations	34,631	34,631	34,631
Adjusted R^2	0.478	0.483	0.833
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A2: Regression Results of Load Factors from Unavailable Products

	Dependent variable:		
	Load Factor from Unavailable Products		
	(1)	(2)	(3)
Distance (1,000 miles)	−0.068***	−0.080***	−0.058***
	(0.004)	(0.004)	(0.003)
Distance^2	0.025***	0.029***	0.019***
	(0.001)	(0.001)	(0.001)
International Hub		0.059***	
		(0.002)	
Origin-Carrier F.E.?	No	No	Yes
Dest-Carrier F.E.?	No	No	Yes
Observations	34,631	34,631	34,631
Adjusted R^2	0.388	0.402	0.671
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A3: First Stage Results of Marginal Cost Estimation

	(1) Baseline	(2) Baseline	(3) Frequency in Demand	(4) Frequency in Demand
Distance	0.050***	0.454***	0.063***	0.511***
	(0.002)	(0.015)	(0.002)	(0.016)
Distance sq.	−0.016***	−0.143***	−0.021***	−0.160***
	(0.001)	(0.005)	(0.001)	(0.005)
Slot Constraint	−0.008***	−0.133***	−0.010***	−0.155***
	(0.001)	(0.005)	(0.001)	(0.006)
Slot Constraint X LCC	0.007***	0.122***	0.010***	0.143***
	(0.002)	(0.015)	(0.002)	(0.017)
Hub	0.012***	0.108***	0.015***	0.121***
	(0.0004)	(0.003)	(0.001)	(0.004)
# NS in Neighbor Mkt on Seg 1	0.003***	0.022***	0.003***	0.024***
	(0.0001)	(0.001)	(0.0002)	(0.001)
LCC in Neighbor Mkt on Seg 1	0.003***	0.035***	0.004***	0.039***
	(0.0003)	(0.003)	(0.0004)	(0.003)
# NS in Neighbor Mkt on Seg 2	0.003***	0.031***	0.004***	0.035***
	(0.0001)	(0.001)	(0.0002)	(0.001)
LCC in Neighbor Mkt on Seg 2	0.001***	0.027***	0.001***	0.032***
	(0.0004)	(0.003)	(0.0004)	(0.003)
Nonstop	−1.071***	−1.229***	−1.087***	−1.216***
	(0.001)	(0.009)	(0.001)	(0.010)
O-D Pair (Airport or City)	Airport	Airport	Airport	Airport
Observations	358,880	358,880	358,880	358,880
Note:			*p<0.1; **p<0.05; ***p<0.01	

Table A4: Airport Codes and Names

Code	Airport Name	City
AGS	Augusta Regional At Bush Field	Bush Field
BDL	Bradley Intl' Airport	Windsor Locks
BHM	Birmingham-Shuttlesworth Intl' Airport	Birmingham
BNA	Nashville Intl' Airport	Nashville
BOS	General Edward Lawrence Logan Intl' Airport	Boston
CAK	Akron Canton Regional Airport	Akron
CHS	Charleston Air Force Base-Intl' Airport	Charleston
CLT	Charlotte Douglas Intl' Airport	Charlotte
CMH	John Glenn Columbus Intl' Airport	Columbus
DAY	James M Cox Dayton Intl' Airport	Dayton
DSM	Des Moines Intl' Airport	Des Moines
FLL	Fort Lauderdale Hollywood Intl' Airport	Fort Lauderdale
GSP	Greenville Spartanburg Intl' Airport	Greenville
IND	Indianapolis Intl' Airport	Indianapolis
JAN	Jackson-Medgar Wiley Evers Intl' Airport	Jackson
JAX	Jacksonville Intl' Airport	Jacksonville
MCI	Kansas City Intl' Airport	Kansas City
MCO	Orlando Intl' Airport	Orlando
MIA	Miami Intl' Airport	Miami
MSP	Minneapolis-St Paul Intl'/Wold-Chamberlain Airport	Minneapolis
MSY	Louis Armstrong New Orleans Intl' Airport	New Orleans
MYR	Myrtle Beach Intl' Airport	Myrtle Beach
OMA	Eppley Airfield	Omaha
PBI	Palm Beach Intl' Airport	West Palm Beach
PNS	Pensacola Regional Airport	Pensacola
PVD	Theodore Francis Green State Airport	Providence
PWM	Portland Intl' Jetport Airport	Portland
ROC	Greater Rochester Intl' Airport	Rochester
RSW	Southwest Florida Intl' Airport	Fort Myers
SAV	Savannah Hilton Head Intl' Airport	Savannah
SRQ	Sarasota Bradenton Intl' Airport	Sarasota
TLH	Tallahassee Regional Airport	Tallahassee
TPA	Tampa Intl' Airport	Tampa
TYS	McGhee Tyson Airport	Knoxville

Table A5: List of New Flight Segments from DCA from the Data

Carrier	New Segment from DCA
Southwest	**FLL**, **TPA**, **BNA**, **MSY**, MCI, **PVD**, **CMH**, **IND**, OMA, MDW
JetBlue	**PBI**, **BDL**, **JAX**, RSW, CHS

This table lists the set of flight segments that Southwest and JetBlue did not provide nonstop services in 2013 (pre-merger) but did in 2016 (post-merger and post-divestiture). A flight segment is in bold face if the model predicts the new flight segment.

Table A6: Logit Model Results of Entry Decision by Southwest and JetBlue

	Dependent variable:
	entered
Variable Profit Change (ΔVP)	1.215**
	(0.477)
Presence (non-DCA)	9.600***
	(2.588)
Constant	−5.982***
	(1.412)
Observations	94
Log Likelihood	−22.790
Akaike Inf. Crit.	51.580
Note:	*p<0.1; **p<0.05; ***p<0.01

Table A7: Difference-in-Differences (DiD) Regression Results of Airplane Size

	Airplane Size			
	(1)	(2)	(3)	(4)
Post X NewAA		0.923	1.427	0.457
		(2.897)	(4.048)	(4.147)
Post	1.341	0.477	−1.021	1.921
	(1.211)	(2.598)	(3.631)	(3.717)
NewAA		−25.027***	−24.200***	−25.872***
		(2.068)	(2.884)	(2.967)
Distance	54.625***	52.166***	51.072***	53.340***
	(1.454)	(1.211)	(1.673)	(1.755)
Constant	43.734***	70.198***	69.997***	70.390***
	(1.208)	(2.073)	(2.898)	(2.968)
DiD Used?	No	Yes	Yes	Yes
Pre-merger Year?	2012-2013	2012-2013	2013	2012
Post-merger Year?	2015-2016	2015-2016	2015	2016
Observations	1,830	2,276	1,147	1,129
Adjusted R^2	0.436	0.517	0.513	0.521
Note:				*p<0.1; **p<0.05; ***p<0.01

B Second Stage Supplements

B.1 Derivations

In this section, I present the technical details of how to derive the matrix-form version of the price FOC (13) from (12). Consider two products j and l offered by carrier f in market m and m', respectively. Then, taking the derivatives of the marginal cost $c_{lm'}$ with respect to p_{jm} from (7) yields the following:

$$\frac{\partial c_{lm'}}{\partial p_{jm}} = \begin{cases} \sum_{s \in \mathcal{S}_{fk}} \gamma_2 \nu \Big(\frac{Q_s}{z_s K_s}\Big)^{\nu} \frac{1}{Q_s} \frac{\partial q_{km}}{\partial p_{jm}} & \text{if } \exists k \text{ s.t. } k \in \mathcal{J}_{mf} \text{ and } k \in \mathcal{J}_f(\mathcal{S}_{fl}) \\ 0 & \text{otherwise} \end{cases} \tag{A.1}$$

Then, the last term of (12), $\sum_{l \in \mathcal{J}_f(\mathcal{S}_f)} \frac{\partial c_{lm'}}{\partial p_{jm}} q_{lm'}$, can be expressed as follows:

$$\begin{aligned} \sum_{l \in \mathcal{J}_f(\mathcal{S}_f)} \frac{\partial c_{lm'}}{\partial p_{jm}} q_{lm'} &= \sum_{k \in \mathcal{J}_{mf}} \sum_{s \in \mathcal{S}_{fk}} \sum_{l \in \mathcal{J}_f(\{s\})} \gamma_2 \nu \Big(\frac{Q_s}{z_s K_s}\Big)^{\nu} \frac{1}{Q_s} \frac{\partial q_{km}}{\partial p_{jm}} q_{lm'} \\ &= \sum_{k \in \mathcal{J}_{mf}} \sum_{s \in \mathcal{S}_{fk}} \gamma_2 \nu \Big(\frac{Q_s}{z_s K_s}\Big)^{\nu} \frac{1}{Q_s} \frac{\partial q_{km}}{\partial p_{jm}} \sum_{l \in \mathcal{J}_f(\{s\})} q_{lm'} \\ &= \sum_{k \in \mathcal{J}_{mf}} \sum_{s \in \mathcal{S}_{fk}} \gamma_2 \nu \Big(\frac{Q_s}{z_s K_s}\Big)^{\nu} \frac{\sum_{l \in \mathcal{J}_f(\{s\})} q_{lm'}}{Q_s} \frac{\partial q_{km}}{\partial p_{jm}} \qquad \text{(A.2)} \\ &= \sum_{k \in \mathcal{J}_{mf}} \sum_{s \in \mathcal{S}_{fk}} \gamma_2 \nu \Big(\frac{Q_s}{z_s K_s}\Big)^{\nu} \Big(\frac{Q_s}{Q_s}\Big) \frac{\partial q_{km}}{\partial p_{jm}}. \qquad \text{(From (8))} \end{aligned}$$

Then,

$$\begin{aligned} \frac{dVP_f}{dp_{jm}} &= q_{jm} + \sum_{k \in \mathcal{J}_{mf}} (p_{km} - c_{km}) \frac{\partial q_{km}}{\partial p_{jm}} - \sum_{l \in \mathcal{J}_f(\mathcal{S}_f)} \frac{\partial c_{lm'}}{\partial p_{jm}} q_{lm'} \\ &= q_{jm} + \sum_{k \in \mathcal{J}_{mf}} \Big(p_{km} - c_{km} - \sum_{s \in \mathcal{S}_{fk}} \gamma_2 \nu \Big(\frac{Q_s}{z_s K_s}\Big)^{\nu}\Big) \frac{\partial q_{km}}{\partial p_{jm}} = 0 \qquad \text{(A.3)} \end{aligned}$$

When we stack all the products in market m in a vector form, we can rewrite expression (A.8) as

$$\mathbf{q_m} + \mathbf{\Omega_m}\Big(\mathbf{p_m} - \mathbf{c_m} - \frac{\mathbf{dc_m}}{\mathbf{dQ}}\Big) = 0, \tag{A.4}$$

where $\mathbf{\Omega_m}$ is the element-wise multiplication of the response matrix defined in (14) and $\frac{\mathbf{dc_m}}{\mathbf{dQ}}$ is a vector with jth element

$$\Big[\frac{\mathbf{dc_m}}{\mathbf{dQ}}\Big]_j = \sum_{s\in\mathcal{S}_{fj}} \gamma_2 \nu \Big(\frac{Q_s}{z_s K_s}\Big)^{\nu}. \tag{A.5}$$

B.2 Missing Products in Data

To precisely estimate the parameters of a load factor and to solve the model using a realistic load factor value, it is crucial to take into account all types of passengers on a segment to. However, not all the products are included in the ticket-level data (DB1B). One example is an international itinerary ticket that contains a domestic flight segment (e.g. a flight ticket from Washington, D.C. (DCA) to Seoul, Korea (ICN) connecting at Detroit, Michigan (DTW), operated by American Airlines). Any tickets including international segments are not in the ticket-level data used in this paper. Another example is a set of tickets to/from small airports that are outside of the top 100 airports based on passenger boardings. Those connecting flights starting from the small airports that contain a DCA flight segment are not in my sample, while passengers from those flights affect the load factor of the DCA segments.

When there are missing products in the data, we can decompose the total number of passengers on flight segment s into two parts—1) a group of passengers for which we have product-level information in the data, and 2) those for whom we do not have product-level information:

$$Q_s = \underbrace{\sum_{j\in\mathcal{J}(\{s\})} q_{jm}}_{\text{Product Available in Data}} + \underbrace{\widetilde{Q_s}}_{\text{Not Available}}. \tag{A.6}$$

Derivations When There Are Missing Products

With this decomposition, we will have slightly different formulas for $\frac{\mathbf{dc_m}}{\mathbf{dQ}}$ in (15) and (19). Basically, we replace $\sum_{l\in\mathcal{J}_f(\{s\})} q_{lm}$ in (A.2) with (A.6) to obtain a new formula:

$$\sum_{l\in\mathcal{J}_f(\mathcal{S}_f)} \frac{\partial c_{lm'}}{\partial p_{jm}} q_{lm'} = \sum_{k\in\mathcal{J}_{mf}} \sum_{s\in\mathcal{S}_{fk}} \gamma_2 \nu \Big(\frac{Q_s}{z_s K_s}\Big)^{\nu} \Big(\frac{Q_s - \widetilde{Q_s}}{Q_s}\Big) \frac{\partial q_{km}}{\partial p_{jm}}. \tag{A.7}$$

Then,

$$\frac{dVP_f}{dp_{jm}} = q_{jm} + \sum_{k\in\mathcal{J}_{mf}} (p_{km} - c_{km}) \frac{\partial q_{km}}{\partial p_{jm}} - \sum_{l\in\mathcal{J}_f(\mathcal{S}_f)} \frac{\partial c_{lm'}}{\partial p_{jm}} q_{lm'}$$

$$= q_{jm} + \sum_{k\in\mathcal{J}_{mf}} (p_{km} - c_{km} - \sum_{s\in\mathcal{S}_{fk}} \gamma_2 \nu \Big(\frac{Q_s}{z_s K_s}\Big)^{\nu} \Big(\frac{Q_s - \widetilde{Q_s}}{Q_s}\Big)) \frac{\partial q_{km}}{\partial p_{jm}} = 0. \tag{A.8}$$

Then, $\frac{\mathbf{dc_m}}{\mathbf{dQ}}$ is a vector with the jth element

$$\Big[\frac{\mathbf{dc_m}}{\mathbf{dQ}}\Big]_j = \sum_{s\in\mathcal{S}_{fj}} \gamma_2 \nu \Big(\frac{Q_s}{z_s K_s}\Big)^{\nu} \Big(\frac{Q_s - \widetilde{Q_s}}{Q_s}\Big), \tag{A.9}$$

where $\widetilde{Q_s}$ denotes the number of passengers from products that are not shown in the ticket-level data. Fortunately, we can observe the total number of passengers at a segment level from the T-100 flight segment-level data (including those not having product information). I leverage this information to obtain the proportion of the load factor on s originating from those products not available in the data, and adjust the load factor when solving the model.

Load Factor Gap Prediction

The fraction of load factor on a new segment coming from unavailable products, $\widetilde{Q_{fs}}$ in (A.6), is predicted. To do so, I first construct the number of passengers on a segment based on T-100 and based on DB1B, respectively. Then, I calculate the load factor gap between the two by their difference divided by the available seats on the segment. I regress this variable on segment specific characteristics such as distance, distance squared, and the international hub dummy. As in the airplane size prediction, I include origin-carrier pairs and destination-carrier pairs fixed effects in the regression. The regression result is reported in Table A2 in the Appendix, and I use coefficients in column (3) to predict the load factor gap on a new segment. This predicted load factor gap will be added whenever the load factor on a flight segment needs to be calculated in the model solving process.

C Counterfactual Supplements

C.1 Marginally Profitable Segments Prediction

I construct the likelihood of new entry in the following way. First, among counterfactual flight segments to/from DCA on which Southwest and JetBlue did not operate any services in the second quarter of 2013 (pre-merger), I identify those new segments on which they initiated nonstop services in the second quarter of 2016 (post-merger), and construct a dummy variable for the entry decisions. Then, I regress the entry decision made by them on the marginal variable profit change due to the newly added segment (ΔVP described in Assumption 1) and the carrier's presence at the non-DCA endpoint. Using the estimated logit coefficient, reported in Table A6, I predict the likelihood of entry for all counterfactual segments of Southwest, JetBlue, Delta, and United. Analogously, a list of marginally profitable routes for NewAA can be obtained by removing each of its pre-merger existing segments from its network, and calculating their likelihood by using the same estimated logit coefficient. Finally, I sort the segments from the most likely marginally profitable to the least, and Table 8 shows the top ten airports that each carrier is likely to enter/exit.

C.2 Post-merger Airplane Size

By using the difference-in-differences methodology, I examine if the merged firm systematically changed its segment-specific airplane size after the merger. To do so, I select the flight segments to/from DCA where the merged firm (treatment group) and other legacy carriers (control group) consistently provided nonstop services from 2012 to 2016. The regression results are shown in Table A7. There is no evidence that the merged firm substantially changed its airplane size at DCA after the merger (column 1) and that ii) the merged firm's segment-level airplane size change is systematically different from that of other legacy carriers' after the merger (column 2). This result is robust under different pre/post merger time horizons (column 3 and 4).

C.3 Heuristic Algorithm

Algorithm 1 shows how the "Heuristic Search" works in my model. Let $\widetilde{\mathcal{S}^0_{f+}}$ and $\widetilde{\mathcal{S}^0_{f-}}$ be the set of marginally profitable segments that carrier f can add to its network and the set of marginally profitable segments that carrier f can remove from its network at the initial period, respectively. $\widetilde{\mathcal{S}^0_{f+}}$ for carriers other than NewAA will be the list of segments

Algorithm 1: Flight Segment Choice Equilibrium Using Heuristic Search

set $n \longleftarrow 0$;
repeat
  $n \longleftarrow n+1$;
  $\Pi_f^{old*} \longleftarrow \Pi_f(\mathcal{S}_f^{n-1}, \mathcal{S}_{-f}^{n-1}), \qquad \forall f \in \mathcal{F}$;
  $\mathcal{S}_f^n \longleftarrow \mathcal{S}_f^{n-1}$, $\widetilde{\mathcal{S}_{f+}^n} \longleftarrow \widetilde{\mathcal{S}_{f+}^{n-1}}$, and $\widetilde{\mathcal{S}_{f-}^n} \longleftarrow \widetilde{\mathcal{S}_{f-}^{n-1}} \qquad \forall f \in \mathcal{F}$;
  for $k \in \{1, 2, ..., K\}$ **do**
    $f \longleftarrow f_k$ where $f_k \in \mathcal{F}$;
    $\Pi_f^{new} \longleftarrow \Pi_f(\mathcal{S}_f^n, \mathcal{S}_{-f}^n)$;
    repeat
      $\Pi_f^{old} \longleftarrow \Pi_f^{new}$;
      Pick $\mathcal{T}_-^* = \arg\max_{\mathcal{T}} \Pi_f(\mathcal{S}_f^n - \mathcal{T}, \mathcal{S}_{-f})$ where $\mathcal{T} \subset \widetilde{\mathcal{S}_{f-}^n}$ and $|\mathcal{T}| \leq 1$;
      Pick $\mathcal{T}_+^* = \arg\max_{\mathcal{T}} \Pi_f(\mathcal{S}_f^n \cup \mathcal{T}, \mathcal{S}_{-f})$ where $\mathcal{T} \subset \widetilde{\mathcal{S}_{f+}^n}$ and $|\mathcal{T}| \leq 1$;
      Denote $\mathcal{T}^*$ one of $(\mathcal{T}_-^*, \mathcal{T}_+^*)$ that gives the higher profit to f ;
      if $\mathcal{T}^* = \mathcal{T}_-^*$ **then**
        $\mathcal{S}_f^n \longleftarrow \mathcal{S}_f^n - \mathcal{T}^*$, $\widetilde{\mathcal{S}_{f-}^n} \longleftarrow \widetilde{\mathcal{S}_{f-}^n} - \mathcal{T}^*$, and $\widetilde{\mathcal{S}_{f+}^n} \longleftarrow \widetilde{\mathcal{S}_{f+}^n} \cup \mathcal{T}^*$;
      else
        $\mathcal{S}_f^n \longleftarrow \mathcal{S}_f^n \cup \mathcal{T}^*$, $\widetilde{\mathcal{S}_{f-}^n} \longleftarrow \widetilde{\mathcal{S}_{f-}^n} \cup \mathcal{T}^*$, and $\widetilde{\mathcal{S}_{f+}^n} \longleftarrow \widetilde{\mathcal{S}_{f+}^n} - \mathcal{T}^*$;
      end
      $\Pi_f^{new} \longleftarrow \Pi_f(\mathcal{S}_f^n, \mathcal{S}_{-f}^n)$;
    until $\Pi_f^{new} > \Pi_f^{old}$;
    Update $\mathcal{S}_f^n$ to $\mathcal{S}_{-g}^n$ of firm g (other than f) ;
    $\Pi_f^{new*} \longleftarrow \Pi_f^{new}$
  end
until $\Pi_f^{new*} > \Pi_f^{old*}, \qquad \forall f \in \mathcal{F}$;

in Table 8, while $\widetilde{\mathcal{S}_{f+}^0}$ for NewAA will be empty set. Analogously, $\widetilde{\mathcal{S}_{f-}^0}$ for NewAA will be the list of segments in Table 8 and for other carriers, will be the empty set. Let $\mathcal{S}_f^0$ be the set of existing segments that f initially operates. Note that $\widetilde{\mathcal{S}_{f-}^0} \subset \mathcal{S}_f^0$. Let $\mathcal{F}$ be the set of carriers and assume that there is a sequence order of carriers $\mathcal{F} = \{f_1, f_2, ..., f_K\}$ where K is the number of carriers.

Given what other carriers choose in the previous period, firm f chooses either one of the two options—adding a new flight segment or removing an existing segment—that gives the firm the highest profit for each turn, and the firm repeats the process until it reaches the point where its profit is no longer increasing. Given firm f's newly updated flight segments, other firms based on order repeat the same thing. After all carriers update their flight segments in this round, the algorithm checks if the profit under the updated segments is greater than the profit in the previous round. The algorithm repeats the same procedure until it reaches the point at which there are no profitable deviations across all carriers.

D Slot Allocation Models

The baseline model that finds the optimal slot allocation can be extended in two directions. One is to change the order in which the choice variables are chosen, and the second is to assume that both demand and marginal cost are affected by a slot allocation.

D.1 Sequential Model

In the baseline model, I assume that a firm *simultaneously* chooses its slot allocation and product prices in the second stage of the model. Alternatively, the firm could choose the slot allocation first, then make product price choices. I call this a sequential model, which actually makes the entry model a three-stage model.

In this sequential model, a slot allocation in the second stage affects product price choices in the third stage. When the carrier excessively allocates its slots to a flight segment, the product costs linked to this slot segment may decline due to the reduced load factor and may give the firm an incentive to reduce prices, which suggests an increase in market competition in markets in which the flight segment is involved. Due to the scarcity of slots, however, the carrier may need to allocate a smaller number of slots to another flight segment. This may lead to a fuller airplane in the segment, and carriers in the markets in which the segment is involved recognize this cost and may set higher prices in those markets. In this model, when setting the optimal slot allocation, carriers recognize that prices are affected by their slot allocation.

D.2 Flight Frequency in Demand

In the baseline model, when a slot allocation changes, a firm's variable profit is affected by this change only through the load factors in the firm's marginal cost. However, it is possible to extend the model by allowing product demand to also be affected by the slot allocation. As the demand estimation result in Table 4 suggests, passengers prefer a flight product that has more daily departure time options. As the number of daily departures of a product increases in the number of landing slots associated with the product, a change in the number of slots assigned to a flight segment can alter both the demand for and the marginal cost of a product.

Equation (A.10) illustrates how adding a slot to flight segment s can affect the variable profit of carrier f, VP_f:

$$\frac{dVP_f}{dK_s} = \underbrace{\sum_j \frac{\partial VP_f}{\partial q_j}\frac{\partial q_j}{\partial K_s}}_{\Delta VP \text{ via Demand}} + \underbrace{\frac{\partial VP_f}{\partial K_s}}_{\Delta VP \text{ via Supply}} . \tag{A.10}$$

An increase in the number of slots assigned to a flight segment not only has a positive and direct impact on the demand for all the products using the slot but also induces a substitution effect between products of the same carrier within a market, as passengers prefer a more- to less-frequent flight product. On the supply side, adding a slot to a flight segment reduces costs for all products using the slot due to the decrease in the load factor via an increase in the number of available seats on the segment.

To illustrate the change in the variable profit from the demand side, Figure 3 presents a simplified carrier airline network. Suppose that carrier f allocates one more slot to the $(\text{DCA}\rightarrow\text{IAH})_f$ flight segment. Then, passengers who fly the DCA-IAH (nonstop) or DCA-IAH-DEN (connecting) products enjoy more departure time options; hence, the slot allocation increases the demand for those products in different markets. However, the marginal increase in demand for the DCA-IAH-DEN product will make some of the passengers in the DCA-DEN market switch to the DCA-IAH-DEN product from other products in the same market.

While this model extension is rich in the sense that it captures both demand and supply responses to a slot allocation, it entails an additional computational burden. In an optimization procedure, a long time is needed to obtain the optimal solution without providing the analytic gradient and/or Hessian of the objective function. The first term in (A.10) makes it challenging to obtain the analytic gradients, as the term is associated with interactions among those products linked through flight segments. I am currently working on this question to increase the computational speed in this model extension.