

Competitive Airline Scheduling under Airport Demand Management Strategies

Vikrant Vaze, Cynthia Barnhart

May 31, 2010

Abstract

Demand often exceeds capacity at the congested airports. Various strategic slot control mechanisms are used to bring demand and supply in balance. Given a slot allocation strategy, the profitability of an airline depends not only on its own schedule but also on competitors' schedules. We propose a game-theoretic model of airline frequency competition under slot constraints. The model is solved to obtain a Nash equilibrium using a successive optimizations approach, wherein individual optimizations are performed using a dynamic programming-based technique. The model predictions are validated against actual frequency data which indicates a close fit to reality. We use the model to evaluate different slot allocation strategies from the perspectives of the various stakeholders. The most significant result of this research shows that a small reduction in total number of slots translates into not only a substantial reduction in airport congestion and delays, but also a considerable improvement in airlines' profitability.

1 Background

Airport congestion is imposing a tremendous cost on the world economy. The US Senate Joint Economic Committee report [30] estimates the total cost of domestic air traffic delays to be around \$41 billion for calendar year 2007, including \$19 billion in additional aircraft operating costs, \$12 billion in passenger delay costs and an estimated \$10 billion in indirect costs of delays to the other industries. The magnitude of these delay costs can be properly grasped by noting that during the same period, the aggregate profits of US domestic airlines were \$4.4 billion [26]. For the year 2007, according to the Bureau of Transportation Statistics [25], delays to around 50% of the delayed flights were categorized as delays caused by the National Aviation System (NAS). Weather and volume were the top two causes of NAS delays, together responsible for

84.51% of the NAS delays. Delays due to volume are those caused by scheduling more airport operations than the available capacity while the delays due to weather are those caused by airport capacity reductions under adverse weather conditions. Both these types of delays are due to scheduling more operations than the realized capacity. Such mismatches between demand and capacity are a primary cause of flight delays in United States.

These delays are disproportionately distributed across airports and metropolitan areas in the country. Congestion at a few major airports is responsible for a large proportion of overall delays. An analysis of air traffic patterns and delays by the Brookings Institution [35] suggests that almost 65% of the delayed flight arrivals are concentrated in the 25 largest metropolitan areas. Moreover, operations across an airline's network are interrelated due to linkages in aircraft, crew and passenger movements. Therefore, delays originating at these major airports propagate across the airline networks causing system-wide disruptive impacts. In the summer of 2007, according to the New York Aviation Rulemaking Committee report [3], 75% of the nationwide flight delays originated from the airports in the New York city area. This suggests that mitigation of demand-capacity imbalance at a handful of congested airports should yield system-wide benefits in terms of delay alleviation.

1.1 Demand Management

Increasing the capacity and decreasing the demand are the two natural ways of bringing the demand-capacity mismatch into balance. Capacity enhancement measures such as building new airports, construction of new runways etc. are investment intensive, require long time horizons, and may not be feasible in several situations due to geographic, environmental, socio-economic and political issues associated with such large projects. On the other hand, demand management strategies such as administrative slot controls, market-based mechanisms or any combinations thereof, have the potential to restore the demand-capacity balance over a medium to short time horizon with comparatively little investment. Demand management strategies refer to any administrative or economic policies and regulations that restrict airport access to users. All the demand management strategies proposed in the literature and practiced in reality can be broadly categorized as administrative controls and market-based mechanisms, although various hybrid schemes have also been proposed.

Administrative Controls: Until very recently, four major airports in United States namely, Laguardia and John F. Kennedy airports at New York, O'Hare airport at Chicago and Reagan airport at Washington D.C., had administrative controls enforced on the number of flight operations. Outside of the US, administrative controls are commonplace at busy airports. Several major airports in Europe and Asia are 'schedule-coordinated', where a central coordinator allocates the airport slots to airlines based on a set of pre-determined rules. Under the current practices, both in and outside of the US, the criteria governing the slot allocation

process are typically based on historical precedents and use-it-or-lose-it rules. An airline is entitled to retain a slot that was allocated to it in the previous year, contingent on the fact that the slot was utilized for at least a certain minimum fraction of time over the previous year. On the other hand, an airline failing to utilize a slot frequently enough is in danger of losing it. One fundamental problem with the current administrative slot allocation procedures is that they are economically inefficient because they create barriers to entry by new carriers [22] and encourage airlines to over-schedule in order to avoid losing the slots [19]. Another fundamental problem, as pointed out by Ball et al. [5], is the implicit need to make a tradeoff between delays and resource utilization. Current approaches require ascertaining the 'declared' capacity of an airport beforehand even though the actual capacity on the day of operations is a function of prevalent weather conditions. Declaring too large a value for capacity poses the danger of large delays under bad weather situations and declaring too low a value leads to wastage of resources under good weather conditions. Declared capacity, i.e. the total number of slots to be allocated per time period, ultimately determines the congestion and delays at an airport.

Congestion Pricing: Researchers have shown that market-based mechanisms, if implemented properly, result in efficient allocation of airport resources. Congestion pricing and slot auctions are two of the most popular market mechanisms proposed in the literature. Classical studies such as Vickrey [37] and Carlin and Park [13] proposed congestion pricing based on marginal cost of delays. Such pricing schemes, in theory, maximize the social welfare through optimal allocation of public resources. Under congestion pricing, the total cost to the user includes the delay cost as well as the congestion price. The notion of equilibrium congestion prices relies on the existence of a demand function, i.e. an expression that gives the aggregate demand for airport resources as a function of total cost to the user. Morrison [23] and Daniel [15] performed numerical experiments to calculate the equilibrium congestion prices under some specific assumptions about the underlying demand function. Carlin and Park acknowledged the problems in estimating demand as a function of congestion prices with any level of reliability.

The aggregate demand for slots at an airport is the sum of the number of slots demanded by each airline. Assuming profit maximizing airlines, the number of slots demanded by an airline can be obtained by equating the incremental profitability of the last slot to the congestion price per slot. In reality, among other factors, the profitability of an airline depends on its own schedule as well as on competitor schedules. It is easy to see that the marginal value of having an extra flight in a particular market largely depends on the number of additional passengers that the airline will be able to carry because of the additional flight, which in turn depends on the schedule of flights offered by the competitor airlines in the same market. So given a set of congestion prices, the total demand for slots should reflect these competitive interactions. Some recent congestion pricing studies, including Pels and Verhoef [28] and Brueckner [12], have modeled competitive effects through Cournot [14] type models of firm competition. However, these models do not incorporate the inverse dependence of one airline's market share on

competitor airlines' frequencies, which is a critical component of such competitive interactions.

Slot Auctions: The idea of airport slot auctions was first proposed by Grether et al. [17]. Rassenti, Smith and Bulfin [29] showed how combinatorial auction design is suitable for airport slot auctions and highlighted the associated efficiency gains through experiments. Since then, several researchers ([7], [6], [22], [19], to name a few) have shown the advantages of slot auctions. The reader is referred to Ball, Donohue and Hoffman [6] and Harsha [19] for a detailed accounts of commonly raised concerns regarding slot auctions and ways of addressing them. In spite of many attractive properties of the auctioning mechanism, an auction by itself does not alleviate airport congestion, but rather allocates a fixed set of resources in a more efficient way. So, to that extent, auctions are similar to administrative controls, as they too pose an implicit need to make a tradeoff between delays and resource utilization.

Once the number of slots to be allocated is determined through some procedure, in theory, slot auctions should maximize the social welfare by allocating the slots to those who value them the most. But the determination of the actual value of a package of slots to an airline is a complicated problem. Harsha [19] proposed a valuation model for estimating the value of a package of slots. However, the formulation does not incorporate any effects of airline competition.

In summary, the problem of managing demand at an airport can be broken down into two types of decisions, which can be taken either sequentially, such as in an auction or administrative mechanism, or simultaneously, such as in a congestion pricing mechanism. An auction or any administrative rule-based mechanism for slot allocation to individual airlines must be preceded by some process that determines the total number of slots per time period. It is this previous step that primarily determines the congestion level. Existing literature has typically focused on the second step and the first step has not received much attention. Furthermore, much of the discussion of the second step excludes any effects of frequency competition. Although congestion pricing tackles both these decisions simultaneously and hence implicitly handles the question of deciding the total number of slots per time period, existing literature on congestion pricing does not incorporate important elements of frequency competition. In this research, we propose a framework for assessing different demand management strategies while explicitly modeling the effects of frequency competition.

1.2 Airline Frequency Competition

Since the deregulation of US domestic airline business in the 1970's, apart from fare, service frequency has become the most important competitive weapon at an airline's disposal. Frequency planning is the part of the airline schedule development process that involves decisions about the number of flights to be operated on each route. By providing more frequency on a route, an airline attracts more passengers. Given an estimate of total demand on a route, the mar-

ket share of each airline depends on its own frequency as well as on the competitor frequency. Market share can be modeled according to the so-called S-curve or sigmoidal relationship between the market share and frequency share, which is a widely accepted notion in the airline industry [24][10]. Empirical evidence of the relationship was documented in some early studies and regression analysis was used to estimate the model parameters [33][34][31]. Over the years, there have been several references to the S-curve including Kahn [21] and Baseler [8]. In a recent study, Wei and Hansen [38] provide further statistical support for the S-curve, based on a nested Logit model for non-stop duopoly markets. The most commonly used mathematical expression for the S-curve relationship [31][10] is given by,

$$MS_i = \frac{(FS_i)^\alpha}{\sum_{j=1}^n (FS_j)^\alpha}. \quad (1)$$

Parameter α is such that $\alpha > 1$. MS_i is the market share of airline i , FS_i is the frequency share of airline i and n is the number of competing airlines.

Despite the continuing interest in frequency competition based on the S-curve phenomenon, literature on the game theoretic aspects of such competition is limited. Hansen [18] analyzed frequency competition in a hub-dominated environment using a strategic form game model. Dobson and Lederer [16] modeled schedule and fare competition as a strategic form game. Adler [1] used an extensive form game model to analyze airlines competing on fare, frequency and aircraft sizes. Each of these three studies adopted a successive optimizations approach to solve for a Nash equilibrium. However, none of these studies assess the impact of starting point on the equilibrium being reached through successive optimizations approach. Also, none of them provide validation of the equilibrium predictions against empirical data. In this research we address both of these limitations. Most of the previous studies involving game theoretic analysis of frequency competition, such as Adler [1], Pels et. al [27], Hansen [18], Wei and Hansen [39], Aguirregabiria and Ho [4], Dobson and Lederer [16], Hong and Harker [20], model market share using Logit or nested Logit type models, with utility typically being an affine function of the inverse of frequency. Such relationships can be considerably different from the S-shaped relationship between market share and frequency share, depending on the exact values of utility parameters. In this research, we use one of the most popular characterizations of the S-curve model.

Furthermore, in most of the previous research, scheduling decisions on one segment are not constrained by the schedule on other segments. This is a good approximation for a situation where an airport is not congested and takeoff and landing slots are freely available. But some congested US airports and several major airports in Europe and Asia are slot constrained. With projected passenger demand in the US expected to outpace the development of new airport capacity, there is a possibility of many more airports in the US employing some form of

demand management in the future. To the best of the authors' knowledge, no previous study has incorporated slot constraints into airline competition models.

1.3 Contributions

The main contributions of this paper fall into four categories. First, we propose a game-theoretic model of frequency competition as an evaluation methodology for slot allocation strategies. Second, we provide a solution algorithm with good computational performance for solving the problem to equilibrium. Third, we provide justification of the credibility of a Nash equilibrium solution concept in two different ways, through empirical validation of the model outcome and through convergence properties of the learning dynamics for non-equilibrium situations. Finally, under simple slot allocation schemes, we evaluate the system performance from the perspectives of various stakeholders, including passengers, airlines and airport operators, and provide insights to guide the demand management policy decisions.

Market-based mechanisms lead to socially efficient resource allocation. But the problems such as calculating the equilibrium congestion prices or ascertaining a set of activity rules for efficient auction design are computationally challenging, even without considering any competitive interactions among the carriers. Therefore, we approach the problem in a different way. Rather than integrating schedule competition into the slot allocations problem, given a slot allocation we provide a framework for predicting the airline schedules and estimating the impact on a variety of stakeholders.

Airline planning process involves a large number of decisions variables. Considerations such as the network effects and demand uncertainty introduce further complications in the process. More tactical decisions such as pricing and revenue management often interact with these planning decisions and hence should be considered in evaluating an airline's response to any slot allocation strategy. Therefore, any mathematical model of airline decisions will involve substantial simplifications and approximations of reality. In this paper, we present the models of airline competition along with a brief discussion of the underlying assumptions and the extent of their validity. After presenting the numerical results, we analyze and estimate the direction and magnitude of the impacts of the main assumptions on the results. In section 2, we provide details of the game-theoretic model of frequency competition under slot constraints. In section 3, we provide an efficient algorithm for equilibrium computation. In section 4, we provide empirical and learning-based justification of the Nash equilibrium outcome. Finally, in section 5, we consider two different slot allocation schemes and evaluate their performance based on multiple criteria. In section 6, we conclude with summary and discussion of the main results.

2 Model

In this section, we describe the relevant notation and formulate the model. In sections 2.1 and 2.2, we will present two important extensions.

We will first formulate the frequency planning problem as an optimization problem from a single airline's point of view. Let us consider an airline a . Consider an airport which is slot constrained, i.e. the number of flights arriving at and departing from that airport is restricted by the slot availability. A slot available to an airline can be used for a flight to or from any other airport, but the total number of slots available to each airline is restricted. An airline's flight frequencies in either direction on a nonstop segment are typically equal or close to each other. In this model, we will only consider the number of flight departures from a slot constrained airport. We will assume that the airports at the other end are not slot constrained. This assumption is quite accurate in the US context, where only a handful of airports are slot constrained. The timing of a slot is also an important aspect of its attractiveness from an airline's point of view. In our model, we focus only on the daily allocation of slots while ignoring the time-of-the-day aspects. To begin with, we will consider frequency planning decisions while assuming that the aircraft sizes remain constant for each segment. We will analyze the impact of this assumption in section 5.2. First we propose a multi-player model of frequency competition where each airline's decision problem is represented as an optimization problem. From here onwards, this model will be referred to as the *basic model*. In this *basic model*, the only decision variable is the number of nonstop flights to be operated on each segment by airline a . This *basic model* is applicable for situations where the fares and other factors are similar among the competing airlines and the main differentiating factor between different airlines is the service frequency. We will relax this assumption in model extension 1 proposed in section 2.1.

A segment is defined as an origin and destination pair for non-stop flights. Let S_a be the set of potential segments with origin at the slot constrained airport. Let p_{as} be the average fare charged by airline a on segment s . Let Q_{as} be the number of passengers carried by airline a on segment s . In general, a passenger may travel on more than one segment to go from his origin to destination, which in some cases involves getting transferred between flights at an intermediate airport. However, we will assume segment-based demand i.e. a passenger traveling on two different segments will be considered as a part of the demand on each segment. This assumption seems quite reasonable for the airports in New York area where nearly 75% of the passengers are non-stop, but not very accurate for major transfer hubs such as the Chicago O'Hare airport. We will analyze the extent of impact of this assumption in section 5.4. Let the total passenger demand on segment s be M_s . C_{as} is the operating cost per flight for airline a on segment s . S_{as} is the seating capacity of each flight of airline a on segment s . Let α_s be the exponent in the S-curve relationship between the market share and the frequency share on the nonstop segment s . The value of α_s depends on the market's characteristics such as long-

haul/short-haul, proportion of business/leisure passengers, etc. The vector of decision variables, $[f_{as}]_{s \in \mathbb{S}_a}$, is the set of flight frequencies for airline a on each of the segments. Because the origin airport is slot constrained, the maximum number of flights that can be scheduled by airline a is restricted to U_a . Often, under the current set of administrative policies based on use-it-or-lose-it type rules, there are restrictions on the minimum number of slots that must be utilized by an airline in order to avoid losing slots for the next year. So there may be a lower limit on the number of slots that must be used. Let L_a be the minimum number of slots that must be utilized by airline a . Let \mathbb{A} be the set of all airlines and let \mathbb{A}_s be the set of airlines operating flights on segment s . As per the S-curve relationship, the market share of airline a on nonstop segment s is $\frac{(f_{as})^{\alpha_s}}{\sum_{a' \in \mathbb{A}_s} (f_{a's})^{\alpha_s}}$, which provides an upper bound on the number of passengers for a specific carrier on a specific segment. This restriction is imposed by constraint (3). Obviously, the number of passengers on a segment cannot exceed the number of seats. However, due to demand uncertainty and due to the effects of revenue management, the airlines are rarely able to sell all the seats on an aircraft. Assuming a maximum achievable load factor of LF_{max} , the seating capacity restriction is modeled by constraint (4). We present results assuming 85% as the maximum load factor value. We also test the sensitivity of the results to variations in this value in section 5.1. The objective function to be maximized is the total operating profit, which is total fare revenue minus total flight operating cost. We have assumed average fares and deterministic demand. We will analyze the impact of these two assumptions in section 5.3. The overall optimization model is as follows,

$$\text{maximize } \sum_{s \in \mathbb{S}_a} (p_{as} * Q_{as} - C_{as} * f_{as}) \quad (2)$$

$$\text{subject to: } Q_{as} \leq \frac{(f_{as})^{\alpha_s}}{\sum_{a' \in \mathbb{A}_s} (f_{a's})^{\alpha_s}} M_s \quad \forall s \in \mathbb{S}_a \quad (3)$$

$$Q_{as} \leq LF_{max} * S_{as} * f_{as} \quad \forall s \in \mathbb{S}_a \quad (4)$$

$$\sum_{s \in \mathbb{S}_a} f_{as} \leq U_a \quad (5)$$

$$\sum_{s \in \mathbb{S}_a} f_{as} \geq L_a \quad (6)$$

$$f_{as} \in \mathbb{Z}^+ \quad \forall s \in \mathbb{S}_a. \quad (7)$$

The market share available to each airline depends on the frequency of other competing airlines in the same market, which in turn are decision variables of those other airlines. Therefore, this is multi-agent model. The optimization problem given by (2) through (7) can only be solved for a given set of values of competitors' frequencies.

Now we propose two extensions to the *basic model*. The first extension is applicable to segments where the competing carriers differ in terms of fare charged or in some other important way. The second extension is applicable to segments on which only one carrier operates nonstop flights.

2.1 Model Extension 1

The *basic model* assumes that the market share depends solely on the frequency share. This assumption is reasonable in many markets where the competitor fares are very close to each other and the competing airlines are similar from the perspectives of the passengers in all other ways. However, for markets where the fares are different, the basic S-curve relationship can be a poor approximation of actual market shares. Consider a market where the competing airlines are differentiated in both fare and frequency. Different types of the passengers would react differently to these attributes. While some passengers value lower fares more, others give more importance to higher frequency and the associated greater flexibility in scheduling their travel. In addition, there could be other airline-specific factors that impact the passenger share. For example, some passengers may have a preference for the big legacy carriers operating wide-body or narrow-body fleet over the regional carriers operating turbo-prop aircraft or small regional jets. To incorporate these effects, we propose an extension of inequality (3). Let there be T types of passengers. Let γ_s^t be the fraction of segment s passengers belonging to type t , such that $\sum_{t=1}^T \gamma_s^t = 1$. Let α_s^t be the frequency elasticity of type t passengers, which serves the same purpose as the exponent of the S-curve in the *basic model*. Let β_s^t be the fare elasticity of type t passengers. Obviously, we expect α_s^t to be non-negative and β_s^t to be non-positive. Let θ_a be the airline specific factor for airline a .

$$Q_{as} \leq \sum_{t=1}^T \frac{\theta_a (f_{as})^{\alpha_s^t} (p_{as})^{\beta_s^t}}{\sum_{a' \in \mathbb{A}_s} \theta_{a'} (f_{a's})^{\alpha_s^t} (p_{a's})^{\beta_s^t}} \gamma_s^t M_s \quad \forall s \in \mathbb{S}_a \quad (8)$$

The market share of each airline is now a function the fares, frequencies and airline specific factors of all competing airlines. This model incorporates the effects of different fares and frequencies on the passenger shares. Also it can model multiple passenger types such as leisure vs. business, by specifying different fare and schedule elasticities for each type of passengers. Also, the remaining airline specific factors are captured through the θ_a parameter.

2.2 Model Extension 2

This model is similar to the *basic model* except that the player decisions are now sequential rather than simultaneous. The idea of modeling the frequency competition as an extensive




Figure 1: Typical shape of segment profit function

form game was proposed by Wei and Hansen [39] where, for contractual or historical reasons, one airline has the privilege of moving first, i.e., deciding the frequency on a segment. The other airline responds upon observing the action by the first player. The *basic model* and the first extension implicitly assumed the existence of at least two competing airlines on a segment. However, frequency decisions in markets with only one existing airline are not completely immune to competition and the incumbent airline must consider the possibility of entry by another competitor while deciding the optimal frequency. Such situations are suitable for being modeled using the idea of Stackelberg equilibrium [32] or a subgame perfect Nash equilibrium of an extensive form game. The incumbent carrier is the Stackelberg leader and the potential entrant is the follower. If a potential new entrant is denoted by a' then the inequality (3) can be extended as,

$$Q_{as} \leq \frac{(f_{as})^{\alpha_s}}{(f_{as})^{\alpha_s} + (f_{a's})^{\alpha_s}} M_s \quad (9)$$

$$f_{a's} = \arg \max_f \left(\min \left(\frac{(f)^{\alpha_s}}{(f_{as})^{\alpha_s} + (f)^{\alpha_s}} * M_s, S_{a's} * f \right) * p_{a's} - C_{a's} * f \right). \quad (10)$$




Figure 2: Typical shape of best response function

3 Solution Algorithm

We will solve the multi-agent game-theoretic model using the Nash equilibrium solution concept. In this section, we describe the solution algorithm used for solving this problem. In section 4, we will provide justification for relying on the Nash equilibrium outcome.

The objective function for each airline is the sum of profits on each segment and the segment frequencies are interrelated through the constraints on the minimum and maximum number of slots. Under the *basic model*, the effect of competitors' frequencies on the profitability of an airline can be completely captured through the notion of effective competitor frequency. Let us define the effective competitor frequency for airline a on segment s as $f_{as}^{eff} = \left(\sum_{a' \in \mathbb{A}_s, a' \neq a} (f_{a's})^{\alpha_s} \right)^{\frac{1}{\alpha_s}}$. So the constraint (3) in the *basic model* can be more succinctly expressed as $Q_{as} \leq \frac{(f_{as})^{\alpha_s}}{(f_{as})^{\alpha_s} + (f_{as}^{eff})^{\alpha_s}} M_s \quad \forall s \in \mathbb{S}_a$. Figure 1 shows the typical form of the segment profit function under the *basic model* for a fixed value of effective competitor frequency, ignoring the slot constraints and the integrality constraints. Under the same assumptions, Figure 2 shows the typical shape of the optimal segment frequency (best response) as a function of effective competitor frequency. The profit function and the best response function gets further complicated by the slot constraints, the integrality constraints, as well as the model extensions 1 and 2. The optimization problem had discrete variables, and as visible from Figure 1, its

continuous relaxation is non-convex. In addition, optimal decisions for each airline depend on the frequency decisions by other airlines. Therefore, the problem of computing an outcome of this multi-agent model can be very challenging. The strategy space for a typical problem size for a major airport is very large with the number of potential candidates for equilibrium solutions of the order of 10^{50} . To solve this problem, we propose a heuristic based on successive optimizations, where individual optimization problems are solved to full optimality using a dynamic programming-based technique.

3.1 Myopic Best Response Algorithm

Let $f_a = [f_{as}]_{s \in \mathcal{S}_a}$ be the vector of frequencies for carrier a . Let $f_{-a} = [f_{a'}]_{a' \in \mathbb{A}, a' \neq a}$ be the vector formed by concatenating the frequency vectors of all competitors of airline a . So any outcome of this problem can be compactly denoted as $f = (f_a, f_{-a})$. Then the myopic best response algorithm (a heuristic) is described as follows,

```

while there exists a carrier  $a$  for whom  $f_a$  is not a best response to  $f_{-a}$  do
     $f'_a \leftarrow$  some best response by  $a$  to  $f_{-a}$ 
     $f \leftarrow (f'_a, f_{-a})$ 
return.

```

This heuristic is based on the idea of myopic best response. Some classes of games have certain desirable properties which make them solvable to equilibrium using an algorithm where each player successively optimizes his own decisions while assuming that the decisions of other players remain constant. Obviously, if such a heuristic converges to some outcome, then it must be a Nash equilibrium. In general, there is no guarantee that it will converge. Further, even if such an algorithm converges to some Nash equilibrium, there is no guarantee that the equilibrium will be unique. We discuss issues regarding the convergence of the myopic best response heuristic and the existence and uniqueness of equilibrium for the game model under consideration in section 4.2.

3.2 Dynamic Programming Formulation

The important building block of the myopic best response algorithm is the calculation of an optimal response of airline a to the competitors' frequencies. Given the frequencies of all competing carriers on all segments, the problem of calculating a best response is an optimization problem. This problem can have a large solution space. For typical problem sizes, the number of discrete solutions in the solution space can be of the order of 10^{10} . As mentioned earlier, this problem is non-convex and discrete. However, this problem has a nice structure. Slot restrictions are

the only coupling constraints and the objective function is additive across individual segments. Therefore, the problem structure is amenable to solution using dynamic programming.

Let $\Pi_s(n)$ denote the profit from operating n flights on segment s . We order the segments and number them from 1 to $|\mathbb{S}_a|$. Segments are considered in order and each segment corresponds to a stage. Each state, (s, n) , corresponds to the combination of the last segment being considered, s , and the cumulative number of flights operated on all the segments considered till then, n . Let $R(s, n)$ be the maximum profit that can be obtained from operating a total of n flights on the first s segments. We initialize $R(0, 0) = 0$ and $R(0, n) = -\infty$ for $n \geq 1$. For $s \geq 1$,

$$R(s, n) = \max_{0 \leq n' \leq n} (R(s-1, n') + \Pi_s(n - n')).$$

The optimal value of total profit for airline a is given by,

$$\max_{L_a \leq n \leq U_a} R(|\mathbb{S}_a|, n).$$

4 Validity of Nash Equilibrium Outcome

Similar to our work, almost all the previous studies on airline competition have used the concept of Nash equilibrium (or one of its refinements) for predicting the outcome of a competitive situation. The traditional explanation for Nash equilibrium is that it results from introspection and detailed analysis by the players assuming that the rules of the game, the rationality of the players, and the profit functions of players are all common knowledge. A Nash equilibrium outcome is attractive mainly because of the fact that unilateral deviation by any of the players does not yield any additional benefits to that player. So given an equilibrium outcome, the players do not have any incentive to deviate from the equilibrium strategies. However, in the absence of any apriori knowledge of an equilibrium outcome, given complicated profit functions such as the ones in this case, it isn't immediately clear why airlines would make the equilibrium decisions. In this section, we substantiate the predictive power of the equilibrium outcome using two different approaches.

4.1 Empirical Validation

Laguardia airport at New York, which has traditionally been one of the few slot constrained airports in the United States, was used for empirical validation of equilibrium frequencies. Flight schedules for US domestic segments are made available by the Bureau of Transportation Statistics [25] for all the certified US carriers that account for at least 1% of the domestic




Figure 3: Empirical validation of frequency predictions

passenger revenue. We compared the equilibrium frequencies predicted by the model against the actual values. For all segments with only one non-stop carrier, the market share function given by model extension 2 was used. The market share function given by model extension 1 was used for segments on which, 1) the competitors' average fares differ by more than 5%, and/or 2) major carriers operating a narrow- or a wide-body fleet compete against regional carriers operating small jets. For all the other segments, the market share function given by inequality (3) in the *basic model* was used. At Lagueardia, the maximum number of slots for each airline is restricted and each airline usually wants to make use of all the slots available to it in order to avoid losing any slots next season. The minimum and maximum numbers of slots available to an airline, i.e. L_a and U_a , are assumed to be equal. So the total number of slots allocated to each airline is fixed. The airline only needs to decide the number of slots to be allocated to flights to each of its destinations. Let f_{as} be the actual frequency of airline a on segment s and \hat{f}_{as} be the equilibrium frequency as predicted by the model. The model ensures that the total frequency for each airline remains constant. Therefore, when the model overestimates the frequency on one segment it necessarily underestimates the frequency on some other segment corresponding to the same carrier. In order to avoid double counting of error, we define a measure of error particularly suitable for such situations. The mean absolute error (MAE) is defined as,

$$MAE = \frac{\sum_{a \in \mathbb{A}} \sum_{s \in \mathbb{S}_a} \max(f_{as}^{\hat{}} - f_{as}, 0)}{\sum_{a \in \mathbb{A}} \sum_{s \in \mathbb{S}_a} f_{as}}.$$

The actual frequency and the frequency predicted by the model for each carrier to each destination is presented in Figure 3. The x-axis lists the combinations of destination airports and airlines. The overall MAE was found to be 7.2%. The model predictions thus match actual frequencies reasonably well.

4.2 Game Dynamics

Airlines typically operate flights on similar sets of segments year after year. The group of competitors on each segment and general properties of markets stay constant over long periods of time. Therefore, the airlines have opportunities to adapt their decisions primarily by fine-tuning the frequency values to optimize their profits. Such adjustments can be captured by modeling the dynamics of the game. In a previous paper, Vaze and Barnhart [36] used a simplified version of the frequency competition model used in this paper and proved the convergence of best response dynamics in the two player case without slot constraints. The key factor responsible for convergence of the myopic best response algorithm was the flat shape of the best response function near equilibrium. In other words, the magnitude of the derivative of the optimal frequency with respect to the effective competitor frequency is very small. Therefore, the best response for a large range of competitor frequency values is very close to the equilibrium frequency, resulting in strong convergence properties of the best response dynamics. The *basic model* of frequency competition used in this research is the same as the model used by Vaze and Barnhart [36], except for the addition of slot constraints and integrality constraints. Though their convergence results are not directly applicable to this complicated model, they provide some intuition.

In this paper, we have used the best response algorithm for computation of an equilibrium. For the results of empirical validation presented in the previous subsection, the vector of actual frequency values was used as the starting point of the best response algorithm and the algorithm was found to converge to equilibrium in 2 iterations per player (i.e. per airline). Let us term this equilibrium solution as the *base equilibrium*. In this section, we present the impact of variation in the starting point on the computed equilibrium prediction. For each starting point, the algorithm was run for at most 10 iterations per player. In most of the following cases, the algorithm converged to an equilibrium and terminated in fewer than 10 iterations. However, in the few cases that the algorithm did not converge within 10 iterations, it was terminated after 10 iterations. Starting from the actual frequency values, we perturbed each dimension of the

Maximum Perturbation	Mean Absolute Error
10%	0.00%
20%	0.79%
30%	1.38%
40%	1.84%
50%	2.45%
60%	3.06%
70%	3.44%
80%	3.80%
90%	4.01%
100%	4.25%

Table 1: Stability of algorithm results to starting point perturbations

frequency vector uniformly between $-x\%$ to $+x\%$ of the original value. For each x value, we drew 1000 samples of starting points. Values presented in Table 1 are the Mean Absolute Error (MAE) values obtained by comparing the solution computed by the best response algorithm to the *base equilibrium*. These results indicate that the model predictions are quite insensitive even to large perturbations in the starting point. The algorithm converges or comes very close to the equilibrium solution within very few iterations, regardless of the starting point. This also suggests that the best response dynamics displays good convergence properties. Therefore, even assuming less than perfectly rational players, an equilibrium outcome can be reached through a simple myopic learning procedure. Empirical validation results, coupled with these desirable convergence properties, make a strong case for using the Nash equilibrium solution concept for predicting the outcome of this airline frequency competition game.

5 Evaluation of Simple Slot Allocation Strategies

In this section, we propose two different strategies for allocating the available slots among different airlines and evaluate the performance of each strategy under the proposed modeling framework.

Proportionate Allocation Scheme: Under the existing administrative controls, airlines often end up receiving a similar number of slots from year to year. Historical precedent is usually used as the main criterion for slot allocation. There is opposition from the established carriers to any significant redistribution of slots. In the spirit of maintaining much of the status quo, the first slot distribution strategy involves proportionate allocation of slots. We vary the

total number of slots at an airport while always distributing them among different carriers in the same ratio as that in reality. For example, if the total number of slots at an airport is reduced from 100 to 80 and if the 100 slots were distributed as 40 and 60 between two carriers, then under our proportionate allocation scheme, the 80 slots will be distributed as 32 and 48 between the same two carriers.

Reward-based Allocation Scheme: While the proportionate allocation scheme is likely to be considered more acceptable by major carriers, it ignores the level of efficiency with which an airline utilizes its slots. Airlines differ, often substantially, in the number of passengers carried per flight or slot. The idea behind the reward-based allocation is to reward those airlines who carry more passengers per slot, due to larger planes and/or higher load factors, and penalize those who carry fewer passengers per slot. Under this scheme, the number of slots allocated to each airline is proportional to the total number of passengers carried by that airline. In the previous example, if the first airline currently carries 140 passengers per slot and the second airline currently carries 120 passengers per slot, then under our reward-based allocation scheme, when the total number of slots is reduced to 80, the first airline will receive 35 and the second airline will receive 45 slots.

We will present results on the impacts of slot reduction assuming that these strategies are implemented as administrative slot controls. Alternatively, these slot allocations can also be thought of as results of some more complicated demand management strategy, such as a market mechanism. If the market mechanism involves monetary payments, the resulting airline profits will have to be adjusted to account for the payments.

5.1 Numerical Results

All the numerical results correspond to Laganardia airport as the slot controlled airport. The data on existing frequencies, fares, aircraft sizes and segment passengers is obtained from the Bureau of Transportation Statistics (BTS) website [25]. The analysis is performed for a weekday in January 2008. Under the proportionate and reward-based slot allocation schemes, for varying numbers of total slots, the profits earned and passengers carried by each airline are computed.

Figures 4 and 5 show the change in total operating profits of all the airlines with slot reductions under the proportionate and reward-based allocation schemes respectively. Figures 6 and 7 show the change in the total number of passengers carried. The total number of passengers carried decreases as the number of slots decreases, but at a much lower rate. For the proportionate allocation scheme, upto a 35% slot reduction, each 1% reduction in slots leads to, on average, just a 0.38% reduction in the total passengers. A 35% reduction in slots leads to approximately 13% reduction in total passengers. Beyond 35%, each 1% reduction in slots leads to slightly over 1% reduction in total passengers. Also, the total operating profit for the proportionate allocation strategy increases with slot reduction percentage upto 35%. Beyond




Figure 4: Total operating profit as a function of slot reductions under a proportionate allocation scheme assuming constant aircraft sizes




Figure 5: Total operating profit as a function of slot reductions under a reward-based allocation scheme assuming constant aircraft sizes




Figure 6: Total number of passengers as a function of slot reductions under a proportionate allocation scheme assuming constant aircraft sizes




Figure 7: Total number of passengers as a function of slot reductions under a reward-based allocation scheme assuming constant aircraft sizes

that point, the operating profit starts decreasing. Very similar patterns are observed for the reward-based allocation strategy. Upto a 40% reduction in slots, each 1% reduction in slots leads to, on average, just a 0.27% reduction in the total passengers. A 40% slot reduction results in less than an 11% reduction in total passengers. However, beyond that point, the rate of reduction in total passengers is close to 1, similar to that in the proportionate reduction case. Similarly, total operating profit increases upto 40% reduction and decreases thereafter. These effects are easy to understand intuitively. Given that aircraft sizes remain constant, the initial reduction in slots results primarily in increases in load factors and hence, under our constant fare assumption, operating costs decrease at a faster rate than the rate of decrease in total revenue. So the operating profit increases. This effect continues until a point where the aircraft size constraint kicks in and reduces the number of passengers almost proportionally to the number of slots. Therefore the operating revenue decreases at almost the same rate as the operating cost decrease, implying that the operating profit decreases. As the total number of slots decreases, the congestion and delays also decrease.

Next, we choose a particular level of slot reduction and evaluate its system-wide impacts based on multiple metrics. As per the airport capacity benchmark report published by FAA [2], the IFR (Instrumental Flight Rules) capacity, i.e. the bad-weather capacity, of Laguardia airport is approximately 87.7% of its optimum capacity. Currently, the number of operations scheduled at Laguardia is close to the good weather (optimum) capacity. We chose to evaluate the case of a 12.3% reduction in slots, which would approximately correspond to scheduling at the IFR capacity instead of the optimum capacity. Although the number of operations scheduled at Laguradia is close to its optimum capacity, the realized capacity drops to the IFR value during bad weather. Estimates of realized capacity values for an entire year, made available from Metron Aviation[®], and actual flight delay data, obtained from the airline on-time performance database available on the BTS website [25], are used for estimating the average delay values. Using the data on realized capacities and flight delays, the average delays to flights landing at Laguardia are calculated for both good and bad weather conditions. The average flight delays under IFR capacity when the number of scheduled operations also equals the IFR value is approximated by the average flight delays under the optimum capacity when the number of operations scheduled also equals the optimum capacity value, which seems reasonable. After the 12.3% slot reduction, the average flight delays under the good weather conditions will be lower than those under good weather conditions without slot reduction. However, in order to be conservative in our delay reduction estimates, we assume that the average delays under good weather conditions remain unchanged upon slot reduction. Finally, the overall average flight delay is calculated as the expected value of delays under good and bad weather capacity.

Table 2 summarizes the impact of slot reduction in terms of flight delays, carrier profits and passengers carried. These results correspond to a 12.3% reduction in slots for both proportionate and reward-based allocation strategies, and the values in parentheses indicate the percentage

Metrics	No Reduction	12.3% Reduction	
		Proportionate	Reward-based
Total Operating Profit	\$1,237,623	\$1,475,217 (19.20%)	\$1,446,520 (16.88%)
Total Passengers	22,184	21,680 (-2.27%)	21,728 (-2.05%)
NAS Delay per Flight	12.74 min	7.52 min (-40.97%)	7.52 min (-40.97%)

Table 2: Effect of a 12.3% slot reduction on system-wide performance metrics

Carrier	No Reduction	12.3% Reduction	
		Proportionate	Reward-based
American Airlines (AA)	\$366,952	\$416,322 (13.45%)	\$406,107 (10.67%)
JetBlue Airways (B6)	\$48,061	\$59,507 (23.82%)	\$59,507 (23.82%)
Continental Airlines (CO)	\$65,996	\$74,466 (12.83%)	\$70,581 (6.95%)
Delta Airlines (DL)	\$196,215	\$252,231 (28.55%)	\$252,900 (28.89%)
AirTran Airways (FL)	\$39,694	\$46,632 (17.48%)	\$48,331 (21.76%)
American Eagle Airlines (MQ)	\$19,831	\$31,318 (57.92%)	\$29,831 (50.43%)
Northwest Airlines (NW)	\$112,578	\$143,084 (27.10%)	\$130,316 (15.76%)
Comair (OH)	-\$1,579	\$39,126 (n.a.)	\$40,582 (n.a.)
United Airlines (UA)	\$208,020	\$224,697 (8.02%)	\$218,922 (5.24%)
US Airways (US)	\$181,855	\$187,834 (3.29%)	\$189,443 (4.17%)

Table 3: Increase in carrier-wise operating profits due to a 12.3% slot reduction

change in each metric. The level of congestion depends on the total number of slots and not on the distribution of these slots among different airlines. Therefore, the delay reduction is the same under both proportionate and reward-based slot allocation strategy. Under either strategy, slot reductions lead to substantial reductions in congestion and delays, considerable increases in operating profits of all carriers, and very small reductions in the total passengers carried.

Table 3 presents the distribution of operating profits across different carriers. Again, the values in parentheses represent the percentage increase in profit. When the total number of slots is reduced under either allocation strategy, the operating profit of each carrier increases compared to that under the existing allocation.

These results are obtained assuming a maximum achievable load factor (LF_{max}) of 85%. Now, we will present results on the sensitivity of the slot reduction impacts to this assumption about the maximum load factor value. Table 4 describes the sensitivity of total profits and

Maximum Load Factor	Increase in Total Profits		Decrease in Total Passengers	
	Proportionate	Reward-based	Proportionate	Reward-based
75%	15.83%	14.33%	2.44%	2.23%
80%	17.39%	17.55%	2.52%	1.94%
85%	19.20%	16.88%	2.27%	2.05%
90%	22.79%	16.44%	0.41%	1.49%
95%	18.90%	17.59%	1.82%	0.94%

Table 4: Sensitivity of slot reduction impacts to the maximum load factor assumption under a 12.3% slot reduction

total number of passengers, to variations in the maximum load factor value. Obviously, the average flight delays do not change, because they depend only on the scheduled number of flight operations. The increase in total operating profit varies between 14.33% and 22.79% and the decrease in number of passengers varies between 0.41% and 2.52%. Due to the integrality constraints on the number of slots, the results don't vary smoothly. But the variation in profit increase and decrease in number of passengers is very small even when the maximum load factor values are varied from 75% to 95%.

5.2 Effect of Aircraft Upgauges

Results in section 5.1 were obtained under the assumption that even when the total number of slots available to an airline is reduced, the airline will continue to operate the same sized aircraft as it did in the absence of slot reduction. This assumption may be realistic for very small reductions in slots, but for significant reductions in slots, it is reasonable to expect that the airlines will tend to operate larger aircraft on some of the segments in order to accommodate more passengers and therefore increase profit. The main problem with modeling aircraft size decisions is that such decisions depend on the fleet availability. In order to estimate the impact of aircraft size upgauges, we extend the model to allow for a limited number of upgauges for each airline. We sort all the available types of aircraft operated out of Lagaardia by any of the airlines in increasing order of seating capacity. We allow a certain maximum percentage of an airline's fleet (operating out of Lagaardia airport) to be upgauged to the next bigger-sized aircraft. This constraint indirectly models the fact that an airline cannot arbitrarily increase aircraft sizes due to fleet availability constraints. For the equilibrium frequency solution under slot reduction situation, the most profitable flight upgauges are performed subject to the limits on maximum allowable upgauge percentage.

Figure 8 describes the impact of limited number of aircraft upgauges on the reductions in




Figure 8: Effect of limited upgauging on total number of passengers under a 12.3% slot reduction

total passengers when the total number of slots is reduced by 12.3%, and the allocation is done based on the proportionate allocation strategy. The maximum allowable upgauge percentage is on the x-axis, which represents the maximum percentage of an airline’s flights that it can upgauge to the next bigger aircraft size. The percentage reductions in the total number of passengers varies from 2.27%, when no upgauges are allowed, to 0.88% when at most 20% upgauges are allowed for each airline. This shows that even with a small fraction of flights upgauging to a larger-sized aircraft, most of the reduction in number of passengers disappears. The remaining reduction in number of passengers is primarily attributable to the fact that there is only a limited number of aircraft sizes available and on some segments, the number of passengers spilled due to a smaller aircraft size is not large enough to justify a profitable upgauge to the next bigger aircraft size.

5.3 Effects of Demand Uncertainty and Revenue Management Practices

The results presented in section 5.1 assumed deterministic passenger demand and the existence of one fixed fare value for each flight. In reality, the passenger demand varies from day to day based on several factors and these variations affect the number of passengers transported. Models of spill estimation typically assume some distribution of demand and calculate the expected value of this distribution truncated by the seating capacity to obtain the expected

Metrics	No Reduction	12.3% Reduction	
		Proportionate	Reward-based
Total Operating Profit	\$1,246,129	\$1,511,805 (21.32%)	\$1,468,370 (17.83%)
Total Passengers	22,347	21,940 (-1.82%)	22,066 (-1.26%)

Table 5: Effect of a 12.3% slot reduction under spill modeling approach

number of passengers carried [11]. So far in our models, to account for demand stochasticity, we put a hard constraint on the maximum number of seats sold on each flight at 85% of the seating capacity. This is a fairly conservative value especially in cases of high demand factors, i.e. high ratios of average passenger demand to seating capacity. Therefore, for high demand factors, we can expect our method to introduce a downward bias in the expected number of passengers carried. Further, the revenue management methods practiced by the airlines affects the average fares of the spilled passengers by ensuring that the spilled passengers are predominantly the low fare passengers. So far we have ignored these effects. Therefore, our method can also be expected to introduce a downward bias in the airline revenue estimates, and hence, in the operating profit estimates. Further, as pointed out by Belobaba and Farkas [11], in addition to impacting the average fares of the spilled passengers, revenue management practices also affect the number of spilled passengers.

Under a slot reduction situation, the passenger carrying capacity is reduced. Revenue management systems used by the airlines are expected to readjust their seat allocation decisions across different fare classes by spilling the low fare passengers, resulting in some increase in average fares. In order to estimate the effect of the demand uncertainty and revenue management systems on the spill and the average fares, we used the approach proposed by Belobaba and Farkas [11] for calculating the expected spill and spill costs using a multiple nested fare class, single booking period model. We will denote our default approach as the deterministic approach and this new approach as the spill modeling approach. For the equilibrium frequency solution with and without slot reductions, the number of passengers carried and the total operating profit of airlines are recalculated using the spill modeling approach. Table 5 shows the impact of a 12.3% slot reduction, using both proportionate and reward-based allocation, in terms of total number of passengers and total operating profits using the spill modeling approach. The numbers in parentheses represent the percentage change as compared to the no reduction case. The results presented in Table 5 correspond to a fare discount ratio of 0.7 for the calculations using the Belobaba and Farkas approach.

The results in Table 5 confirm that our deterministic models introduce a slight downward bias in number of passengers as well as in the operating profits. This can be observed by

comparing these results with those in Table 2. Furthermore, as the demand factor increases, the magnitude of this bias increases, which means that the bias is greater for the slot reduction situation than for the no reduction situation. So our deterministic model, in fact, slightly underestimates the benefits of slot reduction. This is apparent from the greater increase in profits due to slot reduction as reported in Table 5 compared to that in Table 2.

In addition to the effects of demand uncertainty and revenue management, one can consider the possibility of changes in the actual fare values due to reduction in the total number of seats offered on a nonstop segment. However, the authors are not aware of any empirical study on the impacts of seating capacity reduction on the actual fare levels.

All the discussion so far about increases in average fares, either due to preferential spilling of lower fare passengers or due to increases in actual fare levels, assumes that the total seating capacity on a segment is reduced due to slot reduction. However, in light of the possibility of aircraft upgauges as discussed in section 5.2, depending on the extent of upgauging, the decrease in seats due to frequency reduction might be partially or completely compensated by increases in seats per flight. So even the results presented in Table 5 are likely to be slight overestimates of the decrease in passengers and underestimates of the increase in profits.

5.4 Effects of Passenger Connections

Our model of frequency competition assumes segment-based demand. In reality, passengers demand seats on itineraries which may be combinations of two or more flights. Accordingly, the airlines also charge fares on an itinerary basis rather than on a flight basis. So calculation of the fare revenue generated by a segment is not a straightforward process. Belobaba [9] acknowledged that different assumptions made by the airline planners can lead to very different estimates of profitability. The results presented in section 5.1 were calculated under the full fare assumption, i.e. assuming that the entire fare of connecting passengers contributes to the operating profits of each of the segments in the passenger's itinerary. We have considered the effects of slot reduction at a congested airport only on those segments which have one end-point at that airport. Because a very small fraction of passengers (5%) connect at Laganardia, there will be very little double counting of the revenues. So the full fare assumption seems reasonable in that sense. But it fails to capture the fact that not carrying a connecting passenger due to lack of seat availability on only one of the two segments creates the possibility of having an additional nonstop passenger on the other segment, which has seats available. An alternative method for fare revenue calculation is the complete proration approach, in which the fare is completely prorated. Under this assumption, if the airline is not able to carry a connecting passenger, it will only lose the revenue equal to a fraction of the passenger's full fare, which does not make sense if the seat vacated by that passenger on the other segment is not filled by another nonstop or connecting passenger. While neither the full fare assumption nor the

Metrics	No Reduction	12.3% Reduction	
		Proportionate	Reward-based
Total Operating Profit	\$907,248	\$1,067,706 (17.69%)	\$1,121,707 (23.64%)
Total Passengers	22,145	21,116 (-4.65%)	21,751 (-1.78%)

Table 6: Effect of a 12.3% slot reduction under completely prorated fares approach

complete proration assumption is very accurate, these represent the two extremes of possible methods of calculating the segment profitability. Most of the reasonable methods of fare revenue estimation lie somewhere in between these two. Therefore, we evaluated the impact of slot reduction strategies under the complete proration assumption. The results are presented in Table 6. As expected, the profitability values are considerably lower than the corresponding values under the full fare assumption, i.e. those in Table 2. However, the percentage increase in total profits due to slot reduction changes only slightly.

In addition to the fare proration issue, the S-curve based model of market share does not estimate the connecting passenger shares accurately. The S-curve model predicts that when the frequency of one competing carrier decreases compared to the other carrier, a fraction of its passengers will shift to the other carrier. But this does not make sense for connecting passengers if the other carrier does not offer a connecting itinerary to that passenger’s final destination. In fact, with changes in frequency values, the passengers may even shift to traveling through a completely different hub of some other carrier to his final destination. For example, with a reduction in the frequency of American Airlines’ flights from Laguardia (LGA) to Dallas-Fort Worth (DFW) airport, a passenger traveling from Laguardia to Los Angeles (LAX) on American Airlines through its DFW hub may decide to instead travel on Continental Airlines by connecting at Houston (IAH). Therefore, such interactions can get extremely complicated. Of all passengers on flights departing from Laguardia, approximately 75% are nonstop passengers, only 20% are connecting passengers with origin at Laguardia and the remaining 5% connect at Laguardia itself. Of all the connecting passengers originating from Laguardia, nearly 97% are carried by seven major airlines and nearly 89% of them connect through the two biggest hubs of each of these seven airlines. We expect the effects of connecting passengers to be fairly well distributed across these major airlines, without any obvious advantages or disadvantages to any particular carriers. Furthermore, these connecting passenger effects should have a similar impact on results with and without slot reduction. So we expect that the passenger connections would not affect the main results of this study significantly.

6 Summary and Discussion

Any demand management strategy implicitly or explicitly involves deciding the total capacity to be allocated and the distribution of this capacity among different airlines. In this research, we explicitly considered these two stages separately. Although there is extensive literature on airport demand management strategies, none of the previous studies have incorporated critical elements of frequency competition among carriers. To the best of the authors' knowledge, this is the first study that tries to model airline competition under demand management strategies.

We developed a game theoretic model of airline frequency competition based on the S-curve relationship, which is a popular model of market share in the airline literature. Due to the discreteness of the problem and the non-convexity of its continuous relaxation, the optimization problem for each airline is complicated. Furthermore, due to competitive interactions among different players, the problem becomes one of computing a Nash equilibrium. The large size of the solution space makes it very challenging to solve. We used the Nash equilibrium solution concept and proposed an efficient solution algorithm. We justified the predictive power of Nash equilibrium solution concept using empirical validation of the model estimates under existing slot allocation. Irrespective of the starting point, the best response algorithm was found to approach the equilibrium outcome within a very few iterations. This shows that even less than perfectly rational carriers can reach the equilibrium outcome through simple myopic learning dynamics and thus provides further justification of the predictive power of the Nash equilibrium outcome.

We evaluated two simple slot reduction strategies. The results showed that apart from mitigating congestion, small reductions in total allocated capacity can improve the operating profits of carriers. While the two strategies led to considerable differences in the actual profitability increases across individual carriers, the aggregate impacts were similar. Under each strategy, the slot reduction led to substantial increase in the profits of all carriers across the board, and substantial reductions in delays. It also led to a small reduction in the number of passengers carried. However, most of the reduction in total passengers was eliminated when the possibility of a limited amount of aircraft upgauges was introduced. We tested the sensitivity of the results to various assumptions. In most of the cases, the model results varied only slightly. Also in most of the cases, our assumptions were found to be conservative, i.e. relaxing these assumptions is expected to make the slot reduction schemes even more attractive.

Slot reduction is therefore beneficial to the carriers, all of whom experience reductions in delay costs as well as increases in planned operating profits. It also benefits the passengers, almost all of whom get transported to their respective destinations, with significantly lower passenger delays. It is also beneficial to the airport operators because congestion and delays are reduced substantially. From the perspective of the entire system, slot reduction strategies led to almost all passengers being transported with many fewer flights and lower total cost.

Hence, slot reduction strategies are also attractive from the perspective of the overall societal welfare.

References

- [1] N. Adler. Competition in a deregulated air transportation market. *European Journal of Operational Research*, 129(2):337–345, March 2001.
- [2] Federal Aviation Administration. Airport capacity benchmark report 2004. Technical report, Federal Aviation Administration, September 2004.
- [3] Federal Aviation Administration. New york aviation rulemaking committee report. Technical report, Federal Aviation Administration, December 2007.
- [4] V. Aguirregabiria and C. Y. Ho. A dynamic oligopoly game of the US airline industry: Estimation and policy experiments. *Manuscript. University of Toronto*, 2009.
- [5] M. Ball, C. Barnhart, G. Nemhauser, and A. Odoni. Air transportation: Irregular operations and control. *Handbook of OR & MS*, 14, 2007.
- [6] M. Ball, G. Donohue, and K. Hoffman. Auctions for the safe, efficient and equitable allocation of airspace system resources. *Combinatorial Auctions*, 1, 2006.
- [7] M. O. Ball, L. M. Ausubel, F. Berardino, P. Cramton, G. Donohue, M. Hansen, and K. Hoffman. Market-based alternatives for managing congestion at new yorks laguardia airport. In *Proceedings of AirNeth Annual Conference*, 2007.
- [8] R. Baseler. Airline fleet revenue management: Design and implementation. *Handbook of Airline Economics*, 2002.
- [9] P. Belobaba. The airline planning process. In P. Belobaba, A. Odoni, and C. Barnhart, editors, *The Global Airline Industry*, pages 153–181. Wiley, June 2009.
- [10] P. Belobaba. Overview of airline economics, markets and demand. In P. Belobaba, A. Odoni, and C. Barnhart, editors, *The Global Airline Industry*, pages 47–71. Wiley, June 2009.
- [11] P. Belobaba and A. Farkas. Yield management impacts on airline spill estimation. *Transportation Science*, 33(2):217–232, May 1999.

- [12] J. K. Brueckner. Airport congestion when carriers have market power. *The American Economic Review*, 92(5):1357–1375, December 2002. ArticleType: primary_article / Full publication date: Dec., 2002 / Copyright 2002 American Economic Association.
- [13] A. Carlin and R. E. Park. Marginal cost pricing of airport runway capacity. *The American Economic Review*, 60(3):310–319, June 1970.
- [14] A. Cournot. *Recherches sur les principes mathématiques de la thorie des riches, Paris, English translation:(N. Bacon, trans.), Researches into the Mathematical Principles of the Theory of Wealth*. New York: Macmillan & Company, 1897.
- [15] J. I. Daniel. Congestion pricing and capacity of large hub airports: A bottleneck model with stochastic queues. *Econometrica*, 63(2):327–370, March 1995.
- [16] G. Dobson and P. J. Lederer. Airline scheduling and routing in a hub-and-spoke system. *Transportation Science*, 27(3):281297, 1993.
- [17] D. M. Grether, R. M. Isaac, and C. R. Plott. Alternative methods of allocating airport slots: Performance and evaluation. *Pasadena.: Polynomics Research Laboratories, Inc., Tech. Rep*, 1979.
- [18] M. Hansen. Airline competition in a hub-dominated environment: An application of non-cooperative game theory. *Transportation Research Part B: Methodological*, 24(1):27–43, 1990.
- [19] P. Harsha. *Mitigating airport congestion: market mechanisms and airline response models*. PhD thesis, Massachusetts Institute of Technology, Cambridge, Mass, 2008.
- [20] S. Hong and P. T. Harker. Air traffic network equilibrium: Toward frequency, price and slot priority analysis. *Transportation Research Part B: Methodological*, 26(4):307323, 1992.
- [21] A. E. Kahn. Change, challenge, and competition: A review of the airline commission report. *Regulation*, 3:110, 1993.
- [22] Dot Econ Ltd. Auctioning airport slots: A report for HM treasury and the department of the environment, transport and the regions. *Available at hhttp://www.dotecon.comi*, 2001.
- [23] S. A. Morrison. The equity and efficiency of runway pricing. *Journal of Public Economics*, 34(1):45–60, October 1987.

- [24] W. E. O'Connor. *An introduction to airline economics*. Greenwood Publishing Group, 2001.
- [25] Bureau of Transportation Statistics. Airline On-Time statistics and delay causes. http://www.transtats.bts.gov/OT_Delay/OT_DelayCause1.asp?pn=1.
- [26] Bureau of Transportation Statistics. Operating Profit/Loss. http://www.transtats.bts.gov/Data_Elements.aspx?Data=6.
- [27] E. Pels, P. Nijkamp, and P. Rietveld. A note on the optimality of airline networks. *Economics Letters*, 69(3):429–434, December 2000.
- [28] E. Pels and E. T. Verhoef. The economics of airport congestion pricing. *Journal of Urban Economics*, 55(2):257–277, March 2004.
- [29] S. J. Rassenti, V. L. Smith, and R. L. Bulfin. A combinatorial auction mechanism for airport time slot allocation. *The Bell Journal of Economics*, 13(2):402–417, 1982.
- [30] C. E. Schumer and C. B. Maloney. Your flight has been delayed again: flight delays cost passengers, airlines, and the US economy billions. *The US Senate Joint Economic Committee*, 2008.
- [31] R. W. Simpson. *A market share model for U.S. domestic airline competitive markets*. Massachusetts Institute of Technology, Flight Transportation Laboratory, Cambridge, Mass, 1970.
- [32] H. Von Stackelberg. *The theory of the market economy*. William Hodge, 1952.
- [33] N. K. Taneja. *Airline competition analysis*. M.I.T. Flight Transportation Laboratory, Cambridge, 1968.
- [34] N. K. Taneja. *The commercial airline industry*. DC Heath, 1976.
- [35] A. Tomer and R. Puentes. Expect delays: An analysis of air travel trends in the united states. Technical report, Brookings Institution, Washington D.C., October 2009.
- [36] V. Vaze and C. Barnhart. Price of airline frequency competition. Technical Paper. Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, March 2010.
- [37] W. S. Vickrey. Congestion theory and transport investment. *The American Economic Review*, 59(2):251–260, May 1969.

- [38] W. Wei and M. Hansen. Impact of aircraft size and seat availability on airlines' demand and market share in duopoly markets. *Transportation Research Part E: Logistics and Transportation Review*, 41(4):315–327, July 2005.
- [39] W. Wei and M. Hansen. Airlines' competition in aircraft size and service frequency in duopoly markets. *Transportation Research Part E: Logistics and Transportation Review*, 43(4):409–424, July 2007.