EFFICIENT UTILIZATION OF AIRPORT CAPACITY UNDER FREQUENCY COMPETITION

Vikrant Vaze, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Room 1-245, Cambridge, MA 02139 (tel: 617-253-7657, email: vikrantv@mit.edu)

Cynthia Barnhart, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Room 1-206, Cambridge, MA 02139 (tel: 617-253-3815, fax: 617-258-5765, email: cbarnhart@mit.edu)

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 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.

Keywords: airlines, competition, airport, congestion, profitability, game theory, Nash equilibrium

1. BACKGROUND

Airport congestion is imposing a tremendous cost on the world economy. The US Senate Joint Economic Committee report (Schumer and Maloney, 2008) 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 cost, $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 (BTS, 2010b). For the year 2007, according to the Bureau of Transportation Statistics (BTS, 2010a), delays to 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 the 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 (Tomer and Puentes, 2009) 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 (NYARC, 2007), 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 term 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 (DOE, 2001) and encourage airlines to over-schedule in order to avoid losing the slots (Harsha, 2008). Another fundamental problem, as pointed out by Ball et al. (2007b), is the implicit need to make a trade-off 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 (1969) and Carlin and Park (1970) 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 (1987) and Daniel (1995) 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 should reflect these competitive interactions. Some recent congestion pricing studies, including Pels and Verhoef (2004) and Brueckner (2002), have modelled competitive effects through Cournot (1897) type models of firm competition. However, these models do not incorporate the inverse dependence of one airline's passenger demand on competitor airline's frequency, which is a critical component of such competitive interactions.

Slot Auctions: The idea of airport slot auctions was first proposed by Grether et al. (1979). Rassenti, Smith and Bulfin (1982) showed how combinatorial auction design is suitable for airport slot auctions and highlighted the associated efficiency gains through experiments. Since then, several researchers (Ball et al., 2007a; Ball, Donohue and Hoffman, 2006; DOE, 2001; Harsha, 2008; to name a few) have shown the advantages of slot auctions. The reader is referred to Ball, Donohue and Hoffman (2006) and Harsha (2008) for a detailed account 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 trade-off 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 (2008) 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 level of congestion. 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 modelling 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 market share of each airline depends on its own frequency as well as on the competitor frequency. Market share can be modelled 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 (O'Connor, 2001; Belobaba, 2009). Empirical evidence of the relationship was documented in some early studies and regression analysis was used to estimate the model parameters (Taneja, 1968, 1976; Simpson, 1970). Over the years, there have been several references to the S-curve including Kahn (1993) and Baseler (2002). In a recent study, Wei and Hansen (2005) provide 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 (Simpson, 1970; Belobaba, 2009) is given by,

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

Parameter α is such that α > 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 carriers.

Despite the continuing interest in frequency competition based on the S-curve phenomenon, literature on game theoretic aspects of such competition is limited. Hansen (1990) analyzed frequency competition in a hub-dominated environment using a strategic form game model. Dobson and Lederer (1993) modelled schedule and fare competition as a strategic form game. Adler (2001) 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 the successive optimizations approach. Also, none of them provide validation of the equilibrium predictions against actual 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 (2001), Pels, Nijkamp and Rietveld (2000), Hansen (1990), Wei and Hansen (2007), Aguirregabiria and Ho (2009), Dobson and Lederer (1993), Hong and Harker (1992), 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 of a situation where an airport is not congested and the 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 under demand management 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 testing 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 of calculating the equilibrium congestion prices and a set of activity rules for 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 impacts on a variety of stakeholders. 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. We state the assumptions involved in the model and briefly discuss the validity of each of them. In the rest of the paper, we will consider a few relaxations and extensions of the basic model at the appropriate places.

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 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 partially relax this in section 4.1. 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. Q_{as} is 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 travelling on two different segments will be considered as a part of demand on each segment. This assumption seems quite reasonable for the airports in New York area where nearly 80% of the passengers are non-stop, but not very accurate for major transfer hubs such as Chicago O'Hare airport. 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, business/leisure passenger share, etc. The vector of decision variables, $[f_{as}]_{s\in\S_a}$, is the flight frequency for airline a on segment s. Because the origin airport is slot constrained, the total 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 A be the set of all airlines. Let 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 given by $\frac{(f_{as})^{\alpha_s}}{\sum_{a'\in A_s}(f_{a's})^{\alpha_s}}$, which provides an upper bound on the number of passengers for a specific carrier on a specific segment. This restriction is imposed in 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 modelled 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. The objective function to be maximized is the total operating profit, which equals total fare revenue minus total flight operating cost. The overall optimization model is as follows, maximize $\sum_{s\in\S_a}(p_{as} * Q_{as} - C_{as} * f_{as})$ (2)

subject to: $Q_{as} \le \frac{(f_{as})^{\alpha_s}}{\sum_{a'\in A_s}(f_{a's})^{\alpha_s}} * M_s \quad \forall s \in \S_a$ (3)

$Q_{as} \le LF_{max} * S_{as} * f_{as} \quad \forall s \in \S_a$ (4)

$\sum_{s\in\S_a} f_{as} \le U_a$ (5)

$\sum_{s\in\S_a} f_{as} \ge L_a$ (6)

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

The market share available to each airline depends on the frequency of other competing airlines in the same market. But the competitor frequency is a decision variable of other airlines. Therefore, this is multi-agent model. This 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. We propose an extension of inequality (3) where there are 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, which serves the same purpose as the exponent of 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. So the modified inequality (3) is given by,

$$Q_{as} \le \sum_{t=1}^{T} \frac{\theta_a * (f_{as})^{\alpha_s^t} * (p_{as})^{\beta_s^t}}{\sum_{a'\in A_s}\theta_{a'}(f_{a's})^{\alpha_s^t} * (p_{a's})^{\beta_s^t}} * \gamma_s^t * M_s \quad \forall s \in \S_a \tag{8}$$

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

2.2 Model Extension 2

This model is similar to the *basic model* except that the decisions by the two players are now sequential rather than simultaneous. The idea of modelling the frequency competition as an extensive form game was proposed by Wei and Hansen (2007) 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 modelling using the idea of Stackelberg equilibrium (1952) or a sub-game 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 follows,

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

$$f_{a's} = \text{argmax}_f(\min(\frac{(f)^{\alpha_s}}{(f_{as})^{\alpha_s}+(f)^{\alpha_s}} * M_s, S_{a's} * f) * p_{a's} - C_{a's} * f) \tag{10}$$

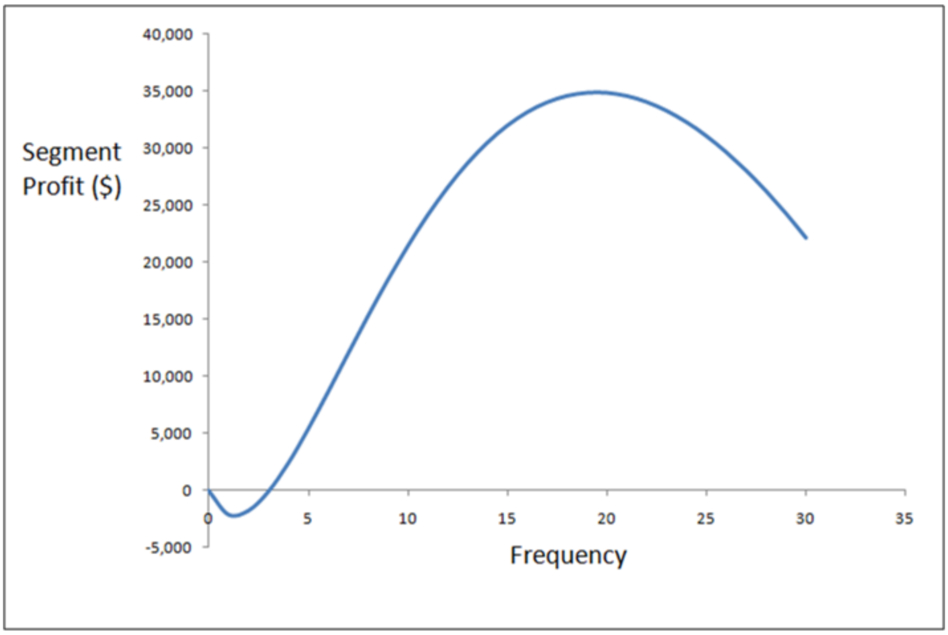

Figure 1 – Typical Shape of Segment Profit Function

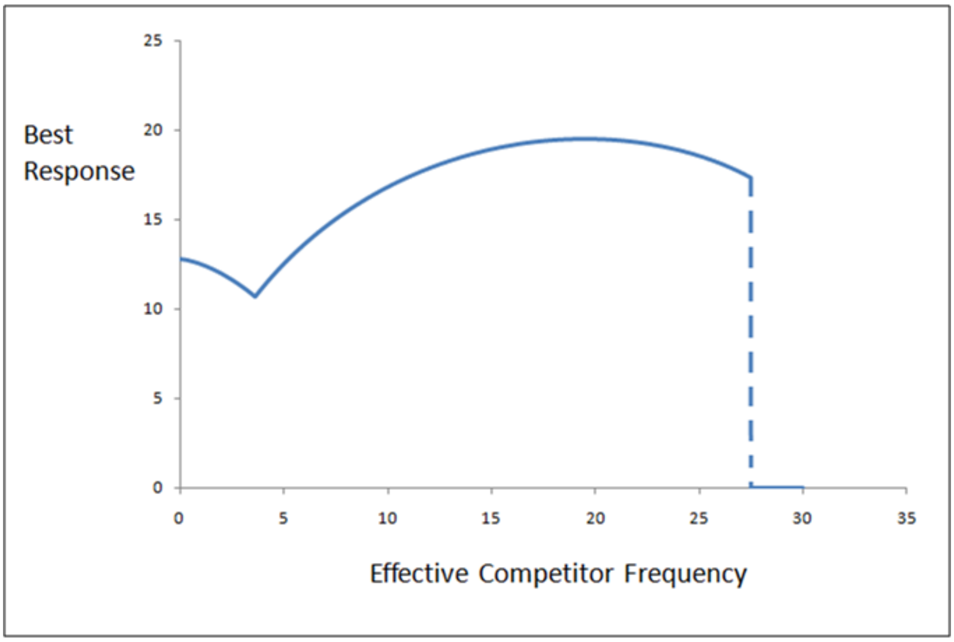

Figure 2 – Typical Shape of Best Response Function

3. SOLUTION ALGORITHM

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 $f_{as}^{eff} = \left(\sum_{a' \in A_s, a' \neq a} (f_{a's})^{\alpha_s^t}\right)^{\frac{1}{\alpha_s^t}}$. So the constraint (3) in the *basic model* can be more succinctly expressed as $Q_{as} \leq \frac{(f_{as})^{\alpha_s}}{(f_{as})^{\alpha_s} + (f_{as}^{eff})^{\alpha_s}} * M_s \quad \forall s \in \S_a$. 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 has 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 the equilibrium outcome 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 being 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 $f_a = [f_{as}]_{s \in \S_a}$ be the vector of frequencies for carrier a. Let $f_{-a} = [f_{a'}]_{a' \in A_s, a' \neq a}$ be the vector formed by concatenating the frequency vectors of all competitors of airline a. So any outcome to this problem can be compactly denoted as $f = (f_a, f_{-a})$. Then the myopic best response algorithm is described as follows:

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

The heuristic is based on the idea of myopic best response. Certain classes of games have attractive properties which make them solvable to equilibrium using an algorithm where each player successively optimizes his own decisions while assuming 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 it converges to some Nash equilibrium, there is no guarantee that that 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 $|\S_a|$. Let $R(s,n)$ be the maximum profit that can be obtained from operating a total of exactly n flights on the first s segments. Segments are considered in order and each segment corresponds to a stage. Each state corresponds to the combination of the last segment being considered and the cumulative number of flights operated on all segments considered till then. 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 $\max_{L_a \leq n \leq U_a} R(|\S_a|, n)$.

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 (2010a) for all the certified US carriers that account for at least 1% of the domestic 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 profit function given by model extension 2 was used. The profit 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 profit function given by equation (3) in the *basic model* was used. At Laguardia, the maximum number of slots for each airline is restricted. Each airline usually wants to make use of all the slots available to it in order to avoid losing the slot next season. Therefore, the minimum and maximum number of slots available to an airline is assumed to be equal. Each airline needs to decide the number of slots to be allocated to flights to each of the potential destinations. Let f_{as} be the actual frequency of airline a on segment s and $\widehat{f_{as}}$ as be the equilibrium frequency 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,

$$MAE = \frac{\sum_{a \in A} \sum_{s \in \S_a} \max(\widehat{f_{as}} - f_{as}, 0)}{\sum_{a \in A} \sum_{s \in \S_a} f_{as}}$$

The actual frequency and frequency predicted by the model for each carrier to each destination is summarized in figure 3. The overall MAE was found to be 7.2%. The model predictions thus match actual frequencies reasonably well.

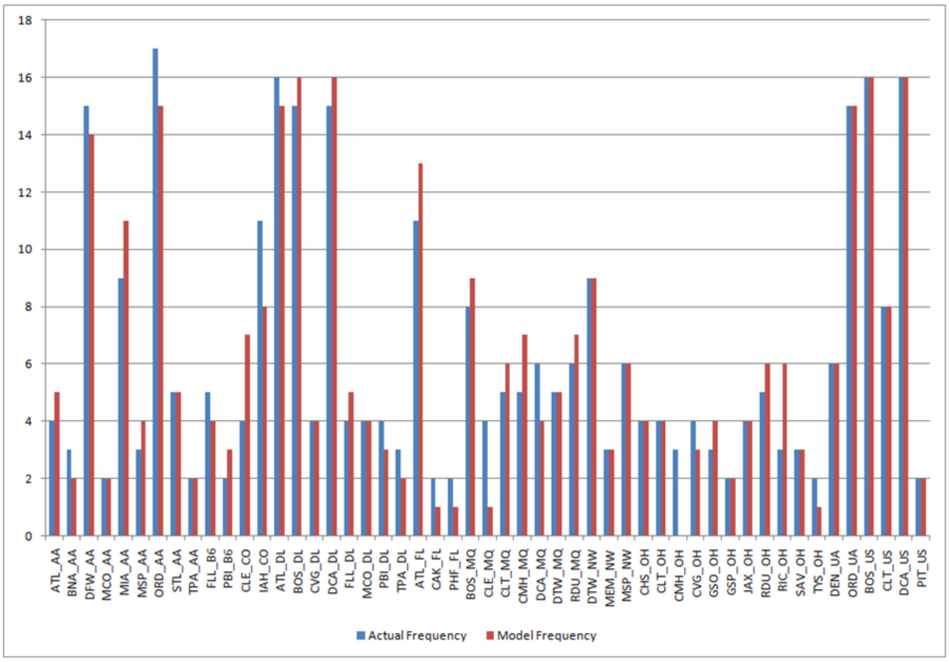

Figure 3 – Empirical Validation of Frequency Predictions

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 modelling the dynamics of the game. In a previous paper, Vaze and Barnhart (2010) 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 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 low. Therefore, the best response to a large range of competitor frequency values is very close to the equilibrium frequency, resulting in strong convergence properties of the best response dynamics to equilibrium. The *basic model* of frequency competition used in this research is the same as the model used by Vaze and Barnhart (2010), 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 cases, the algorithm converged to an equilibrium and terminated in less 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 frequency vector uniformly between -*x*% to +*x*% of the original value. For each *x* value, we drew 1000 samples of starting point. 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 predictions are quite insensitive even to large perturbations to the starting point. Also, 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 dynamic displays good convergence properties. Therefore, even assuming less than perfectly rational players, an equilibrium outcome can be reached through a simple myopic learning procedure.

Table I – Stability of Algorithm Results to Starting Point Perturbations

Maximum Perturbation	MAE
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%

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 modelling 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.

All the numerical results presented in this section correspond to Laguardia airport as the slot controlled airport. All the results excluding section 5.1 assume that the aircraft sizes for each airline on each segment remain unchanged. The data on existing frequencies, fares, aircraft sizes and segment passengers is obtained from the Bureau of Transportation Statistics website (2010a). The analysis is performed for a weekday in January 2008. Under these two allocation schemes, for varying numbers of total slots, the profits earned and passengers carried by each airline are computed using our modelling framework and solution algorithm. The level of congestion depends on the total number of slots. Estimates of realized capacity values for an entire year were made available from Metron Aviation®. Conservative estimates of the delay reductions were obtained from analyzing actual delay data under these realized capacity values.

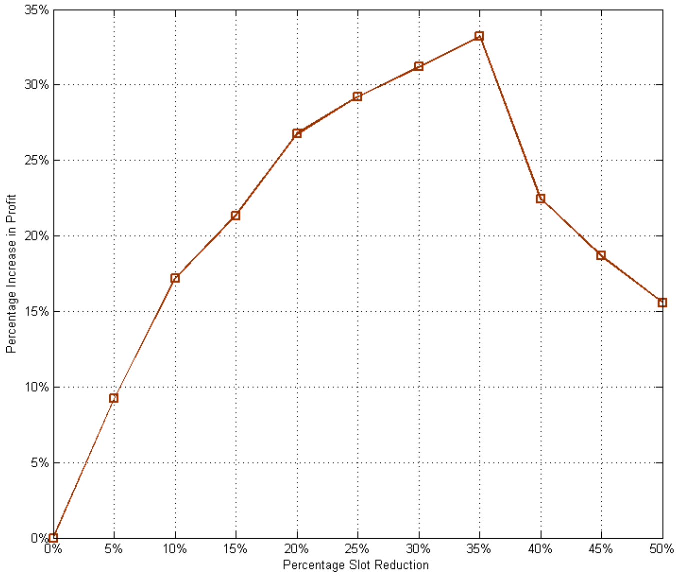

Figure 4 – Total Operating Profit as a Function of Slot Reductions under a Proportionate Allocation Scheme Assuming Constant Aircraft Sizes

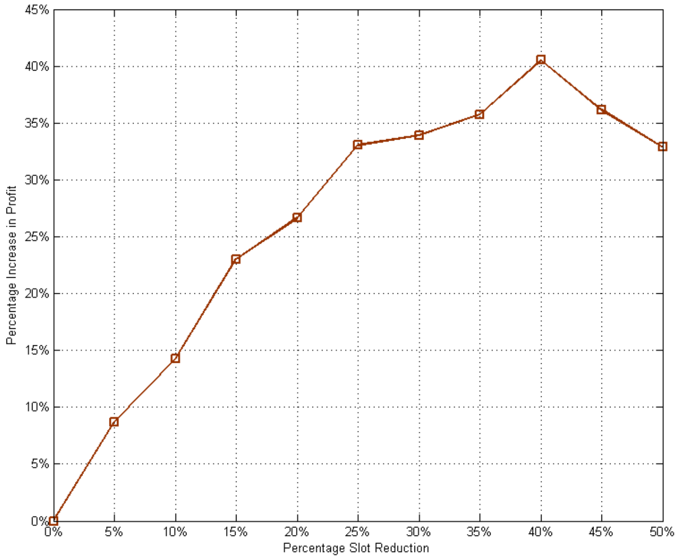

Figure 5 – Total Operating Profit as a Function of Slot Reductions under a Reward-based Allocation Scheme Assuming Constant Aircraft Sizes

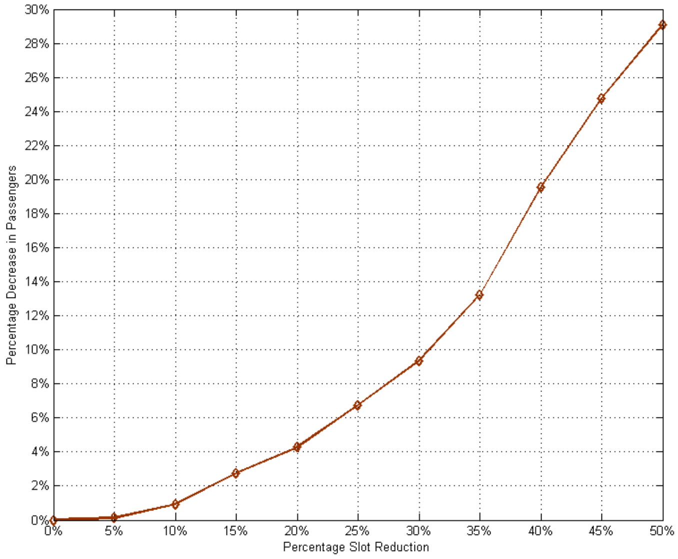

Figure 6 – Total Number of Passengers as a Function of Slot Reductions under a Proportionate Allocation Scheme Assuming Constant Aircraft Sizes

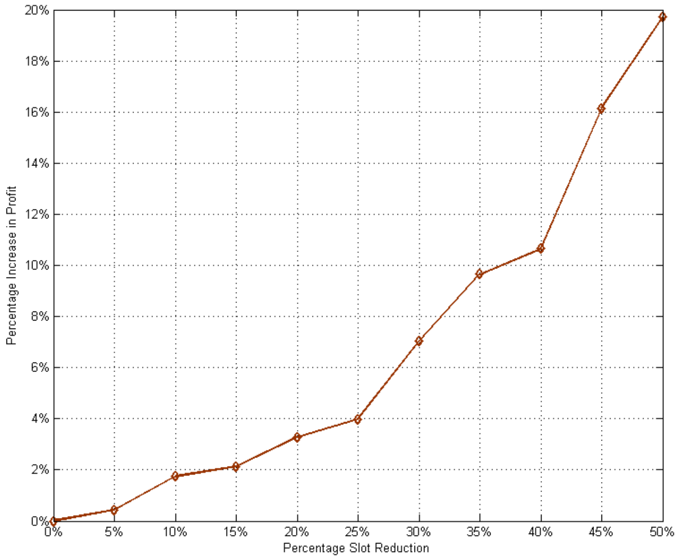

Figure 7 – Total Number of Passengers as a Function of Slot Reductions under a Reward-based Allocation Scheme Assuming Constant Aircraft Sizes

Figures 4 and 5 show the change in total operating profit 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, up to 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 13.2% reduction in total passengers. Beyond 35%, each 1% reduction in slots leads to nearly a 1% reduction in total passengers. Also the total operating profit for the proportionate allocation strategy increases with slot reduction percentage up to 35%. Beyond that point, the operating profit starts decreasing. Very similar patterns are observed for the reward-based allocation strategy. Up to a 40% reduction in slots, each 1% reduction leads to, on average, just a 0.27% reduction in the total passengers. A 40% slot reduction results in less than an 10.7% 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 up to the 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 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, implying that the operating profit decreases. As the total number of slots decreases, the congestion and delays also decrease.

Next, we fix a particular level of slot reduction and evaluate its impact based on multiple criteria. As per the airport capacity benchmark report published by FAA (2004), the IFR (Instrumental Flight Rules) capacity, that is, the bad-weather capacity, at 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. The results in Tables 2 and 3 correspond to a 12.3% reduction in slots, which would approximately correspond to scheduling at the IFR capacity values instead of the optimum capacity values. Table 2 presents the operating profits for each carrier in the case of no slot reduction and in the case of 12.3% slot reduction using both the proportionate allocation strategy and the reward-based allocation strategy. The numbers in parentheses represent the percentage increase in profit. When the total number of slots is reduced under either strategy, the operating profit of each carrier increases compared to that under the existing allocation.

Table 2 – Increase in Operating Profits due to 12.3% Slot Reduction

Carrier	Existing Allocation	IFR Proportionate	IFR 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 (-)	$40,582 (-)
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 summarizes the impact of slot reduction in terms of congestion alleviation, carrier profits and passengers carried. Again, these results correspond to a 12.3% reduction in slots for both proportionate and reward-based allocation strategies, and the numbers in parentheses indicate the percentage change in each metric. 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 number of passengers carried.

Table 3 – Effect of 12.3% Slot Reduction on System-wide Performance Metrics

Strategy	Existing Allocation	IFR Proportionate	IFR 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%)

So far, we have assumed that the maximum load factor (LF_{max}) is 85%, due to the effects of demand uncertainty and revenue management practices. We tested the sensitivity of the system-wide impacts to this assumption about the maximum load factor value. Table 4 and Table 5 describe the sensitivity of total profits and total number of passengers respectively, to variations is maximum load factor value. It can be observed that upon varying the maximum load factor value in the range 75% to 95%, the increase in total operating profit varies between 14.33% and 22.79%, and the decrease in total number of passengers varies between 0.41% and 2.52%.

Table 4 – Increase in Total Profits (under 12.3% Reduction) for Different Maximum Load Factor Values

Maximum Load Factor	Proportionate Reduction	Reward-based Reduction
75%	15.83%	14.33%
80%	17.39%	17.55%
85%	19.20%	16.88%
90%	22.79%	16.44%
95%	18.90%	17.59%

Table 5 – Decrease in Number of Passengers (under 12.3% Reduction) for Different Maximum Load Factor Values

Maximum Load Factor	Proportionate Reduction	Reward-based Reduction
75%	2.44%	2.23%
80%	2.52%	1.94%
85%	2.27%	2.05%
90%	0.41%	1.49%
95%	1.82%	0.94%

5.1 Relaxing the Constant Aircraft Size Assumption

So far we presented results based on 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 before. This might 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 modelling aircraft size decisions is that such decisions depend on the fleet availability. In order to estimate the impact of aircraft upgauges, we extend the model so that the carriers can increase the aircraft size to some extent. We sort all the available types of aircraft operated out of Laguardia by any of the airlines in increasing order of seating capacity. We extend the model to evaluate the impact of allowing a certain maximum percentage of an airline's fleet (operating out of Laguardia airport) to be upgauged to the next bigger-sized aircraft. This constraint indirectly models the fact that, due to fleet availability constraints, an airline cannot arbitrarily increase aircraft sizes.

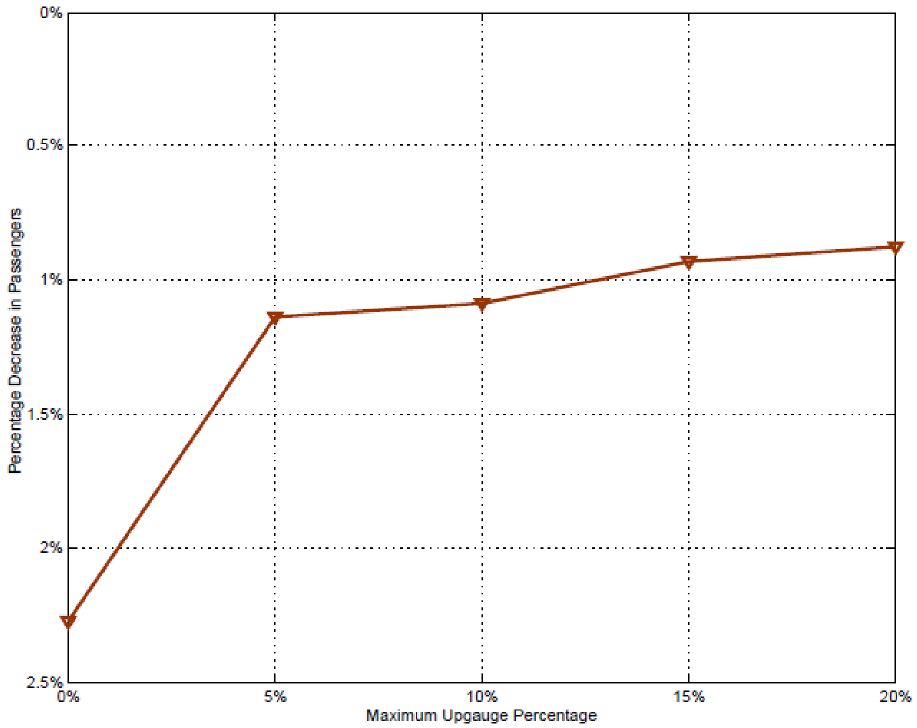

Figure 8 – Effect of Limited Upgauging on Total Number of Passengers

Figure 8 describes the impact of larger aircraft sizes on the reductions in 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 reduction 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. The small remaining decrease in passengers at 20% maximum upgauge level is nothing but an effect of the finite set of aircraft sizes being available. Even this decrease would disappear if we allow for finer discretization of aircraft sizes. These results show that even with a small fraction of flights upgauging to a larger-sized aircraft, most of the reduction in total passengers disappears.

6. SUMMARY AND DISCUSSION

In this research, we develop 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 an empirical validation of the model results under existing slot controls. Irrespective of the starting point, the best response algorithm approaches 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. This provides further justification of the the predictive power of the Nash equilibrium outcome.

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 consider 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.

Given the modelling framework and tools for solving it to equilibrium, we evaluated two simple slot allocation 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 increases in the profits of all carriers across the board, and substantial reductions in flight 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. So, slot reduction is beneficial to the carriers who each experience reductions in delay costs as well as increase in planned operating profit. This benefits passengers, almost all of whom get transported 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, all passengers are transported with many fewer flights and lower total cost. Hence, slot reduction strategies are also attractive from the perspective of overall societal welfare.

Thus, when airline frequency competition is incorporated into the evaluation of demand management strategies, demand management can be shown to be beneficial to the competing carriers. We conclude that simple demand management strategies involving a small reduction in total allocated capacity reduce congestion and are attractive to the various stakeholders and the system as a whole.

REFERENCES

Adler, N. (2001). Competition in a deregulated air transportation market. European Journal of Operational Research, 129(2), 337-345.

Aguirregabiria, V. And C.Y. Ho (2009). A dynamic oligopoly game of the US airline industry: estimation and policy experiments. Manuscript. University of Toronto.

Ball, M., G. Donohue and K. Hoffman (2006). Auctions for the safe, efficient and equitable allocation of airspace system resources. Combinatorial Auction, 1.

Ball, M., L.M. Ausubel, F. Berardino, P. Cramton, G. Donohue, M. Hansen, and K. Hoffman (2007a). Market-based alternatives for managing congestion at New York's Laguardia Airport. In: Proceedings of Airneth Annual Conference.

Ball, M., C. Barnhart, G. Nemhauser and A. Odoni (2007b). Air transportation: irregular operations and control. Handbook of OR & MS, 14.

Baseler, R. (2002). Airline fleet revenue management: design and implementation. Handbook of Airline Economics.

Belobaba, P. (2009). Overview of airline economics, markets and demand. In: P. Belobaba, A. Odoni and C. Barnhart eds., The Global Airline Industry, pp 47-71, Wiley.

Brueckner, J.K. (2002). Airport congestion when carriers have market power. The American Economic Review, 92(5):1357-1375.

Bureau of Transportation Statistics (BTS, 2010a). Airline on-time statistics and delay causes. Available at: http://www.transtats.bts.gov/OT_Delay/OT_DelayCause1.asp. Accessed on: Apr 18, 2010.

Bureau of Transportation Statistics (BTS, 2010b). Operating profit/loss. Available at: http://www.transtats.bts.gov/Data_Elements.aspx?Data=6. Accessed on: Apr 18, 2010.

Carlin, A. And R.E. Park (1970). Marginal cost pricing of airport runway capacity. The American Economic Review, 60(3):310-319.

Cournot, A. (1897). Researches into the mathematical principles of the theory of wealth. Macmillan & Company, New York.

Daniel, J.I. (1995), Congestion pricing and capacity of large hub airports: a bottleneck model with stochastic queues. Econometrica, 63(2):327-370.

Dobson, G. and P.J. Lederer (1993). Airline scheduling and routing in a hub-and-spoke system. Transportation Science, 27(3):281-297.

Dot Econ Ltd. (DOE, 2001). Auctioning airport slots: a report for HM treasury and the department of the environment, transport and the regions. Available at http://dotecon.com/publications/slotauctp.pdf.

Federal Aviation Administration (FAA, 2004). Airport capacity benchmark report 2004. Technical report.

Grether, D.M., R.M. Isaac and C.R. Plott (1979). Alternative methods of allocating airport slots: performance and evaluation. Pasadena Polynomics Research Laboratories, Technical report.

Hansen, M. (1990). Airline competition in a hub-dominated environment: an application of non-cooperative game theory. Transportation Research Part B: Methodological, 24(1):27-43.

Harsha, P. (2008). Mitigating airport congestion: market mechanisms and airline response models. PhD Thesis, Massachusetts Institute of Technology, Cambridge Mass.

Hong, S. And P.T. Harker (1992). Air traffic network equilibrium: toward frequency, price and slot priority analysis. Transportation Research Part B: Methodological, 26(4):307-323.

Kahn, A.E. (1992). Change, challenge and competition: a review of the airline commission report. Regulation, 3:110.

Morrison, S. (1987). The equity and efficiency of runway pricing. Journal of Public Economics, 34(1):45-60.

New York Aviation Rulemaking Committee (NYARC, 2007). New York aviation rulemaking committee report. Technical report.

O'Connor, W.E. (2001). An introduction to airline economics. Greenwood Publishing Group.

Pels, E., P. Nijkamp and P. Rietveld (2000). A note on the optimality of airline networks. Econometric Letters, 69(3):429-434.

Pels, E. And E.T. Verhoef (2004). The economics of airport congestion pricing. Journal of Urban Economics, 55(2):257-277.

Rassenti, S.J., V.L. Smith and R.L. Bulfin (1982). A combinatorial auction mechanism for airport time slot allocation. The Bell Journal of Economics, 13(2):402-417.

Schumer, C.E. and C.E. Maloney (2008). Your flight has been delayed again: flight delays cost passengers, airlines, and the US economy billions. The US Senate Joint Economic Committee Report.

Simpson, R.W. (1970). A market share model for US domestic airline competitive markets, MIT Flight Transportation Laboratory, Cambridge Mass.

Taneja, N.K. (1976). The commercial airline industry. DC Health.

Taneja, N.K. (1968). Airline competition analysis. MIT Flight Transportation Laboratory, Cambridge Mass.

Tomer, A. And R. Puentes (2009). Expect delays: an analysis of air travel trends in the United States. Technical report, Brookings Institution, Washington DC.

Vaze, V. and C. Barnhart (2010). Price of airline frequency competition. Working paper, Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge Mass.

Vickrey, W.S. (1969). Congestion theory and transport investment. The American Economic Review, 59(2):251-260.

Von Stackelberg, H. (1952). The theory of the market economy. William Hodge.

Wei, W. and M. Hansen (2005). 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.

Wei, W. and M. Hansen (2007). Airlines' competition in aircraft size and service frequency in duopoly markets. Transportation Research Part E: Logistics and Transportation Review, 43(4):409-424.