Two-Stage Game Theoretic Modelling of Airline Frequency and Fare Competition

Reed Harder
Thayer School of Engineering
Dartmouth College
reed.haseltine.harder.th@dartmouth.edu

Vikrant Vaze
Thayer School of Engineering
Dartmouth College
vikrant.s.vaze@dartmouth.edu

ABSTRACT

Airlines make decisions about pricing and daily service frequency in a competitive environment. We develop a two-stage game-theoretic model of airline competition, where airlines make frequency decisions during the first stage and fare decisions during the second stage with knowledge of the first stage decisions. We prove that for a simplified two-player form of this game, with assumptions of unrestricted seats-per-flight and only non-nonstop passengers, the first-stage payoff function of each player is concave with respect to that players frequency strategy. With the same assumptions, we also prove that this two-stage game belongs to the class of sub-modular games. Concavity and sub-modularity are shown by numerical experiments to hold for one player, two player, and three player games across a wide range of parameter values, with quadratic functions of player frequencies providing a good approximation ($R^2 > 0.9$) for airline payoffs in all cases. We use solve this model for an 11-airport, four-airline network using the myopic best-response learning heuristic, and the frequency predictions from this solution are validated against actual frequency data from this network. This paper demonstrates that a two-stage frequency-fare game of airline competition can exhibit properties (concavity and sub-modularity) that allow for a computationally tractable equilibrium solution across a wide range of parameter values and a good fit with observed airline frequencies.

Keywords

Aviation, Game-Theoretic Modelling

1. INTRODUCTION

Airlines make capacity and fare decisions in a competitive environment. Capacity decisions, encompassing decisions about frequency of service and seats-per-flight, affect both the operating costs and revenues of airlines. These decisions have significant implications for the performance of the air transportation system as a whole. Over- and under-allocation of airline capacity has been shown to result in billions of dollars in costs to airlines and passengers, wastage of system resources, passenger inconvenience, and environmental damages ([12] and others, refer to [9] for a full review). Airline frequency competition in particular has been shown to be a major driver of increased airport congestion [15]. Capacity and fare decisions of different airlines are interdependent, both serving as tools in an airlines competitive arsenal. The interdependency of different airlines decisions motivates a game theoretic approach to modeling their decision process. A validated game theoretic model of airline decision-making could support forecasting of airline behavior in different scenarios, and provide insight into the impacts of different policy rules on air transportation services. Past studies have developed single-stage game theoretic models considering frequency competition (e.g., [7]), and capacity and fare competition (e.g., [4]). In reality, however, capacity and fare decisions are made sequentially by different departments within an airline. Capacity (especially frequency) decisions are typically made months in advance of flight departure, with only an approximate knowledge of what fares will be, while fare decisions are made weeks to minutes ahead of flight departure. Several studies have stressed the need to develop two-stage game theoretic models to account for the sequential nature of these decisions ([6] and [8], among others; see [9] for a full review), but there are very few studies that bridge analytical, computational, and empirical results for such models. We develop a two-stage frequency and fare competition model, and solve for a subgame perfect Nash equilibrium. We prove various analytical properties; demonstrate its tractability across a wide range of assumptions, scenarios, and parameter values; and validate its predictions against observed airline behavior.

2. MODEL AND ANALYTICAL RESULTS

We focus on the frequency component of capacity decisions, while holding seats-per-flight constant. Seats-per-flight decisions, while significant, have limited effect on passengers' itinerary choice, and show significantly lower variability than frequency decisions across time and across flight segments within an airline [9]. The payoff function for each airline is given as the difference between revenue and operating costs. For an airline a in market m (an origin-destination pair), revenue is computed as

$$Rev_{a,m} = \min\left(M_m \mathrm{MS}_{a,m}, f_{a,m} s_{a,m}\right) p_{a,m}$$

where M_m is the market size, $s_{a,m}$ is the seats-per-flight, $p_{a,m}$ is the fare, $f_{a,m}$ is the daily flight frequency and $\mathrm{MS}_{a,m}$ is the airline's market share. We use a multinomial logit model of market share, an approach widely used in literature. Logit passenger utility is given by a linear combination of fare and a logarithmic transformation of frequency. Operating costs are modeled as a linear function of frequency on a given segment, such that $Cost_{a,m} = c_{a,m} f_{a,m}$, where $c_{a,m}$ is the operating cost per flight. Following the s-curve model of airline frequency competition, we can model market share for carrier a in market m by

$$\mathrm{MS}_{a,m} = \frac{e^{\alpha_m \ln(f_{a,m}) - \beta_m p_{a,m}}}{N_m + \sum_{i \in A_m} e^{\alpha_m \ln(f_{i,m}) - \beta_m p_{i,m}}} \qquad (1)$$

Here, A_m is the set of airlines in market m, α_m and β_m are the utility parameters, and N_m is the utility of the no-fly alternative. In this model, in the absence of a no fly option, $\alpha = 1$, and all else equal, frequency share determines market share.

Market share can also follow the schedule delay model, as discussed in [8], which takes into account the discrepancy between available flight times and the flight times desired by passengers. In this case, the market share is given by

$$\mathrm{MS}_{a,m} = \frac{e^{-\varphi f_{a,m}{}^{-r} - \beta p_{a,m}}}{N_m + \sum_{i \in A_m} e^{-\varphi f_{i,m}{}^{-r} - \beta p_{i,m}}} \qquad (2)$$

Here, φ and r are frequency parameters. Thus, with either the s-curve or schedule delay market share model, the payoff function of airline a operating in a set of markets K_a is given by

$$\pi_a = \sum_{m \in K_a} \min\left(M_m \mathrm{MS}_{a,m},\ f_{a,m} s_{a,m}\right) p_{a,m} - c_{a,m} f_{a,m}$$

Frequency decisions, $f_{a,m}$, are made in the first stage of the game, while fare decisions, $p_{a,m}$, are made in the second stage. We begin our analysis with a simplified version of this game: two airlines competing in a single market, with no connecting passengers, infinite seating capacity, and the absence of a no-fly alternative. Under these assumptions, for either the s-curve market share model (1), or the schedule delay market share model with some parametric assumptions (2) we are able to prove the following propositions:

PROPOSITION 1. *The second-stage fare game always has unique pure strategy Nash equilibrium.*

PROPOSITION 2. *In the first-stage frequency game, each airline's payoff π_i for $i \in \{1,\ 2\}$ is concave in airline i's own strategy, across plausible parameter ranges.*[1] *That is,* $\frac{\partial^2 \pi_i}{\partial f_i{}^2} < 0$ *for* $i \in \{1,\ 2\}$.

PROPOSITION 3. *In the first-stage frequency game, each airline's payoff π_i for $i \in \{1,\ 2\}$ is a submodular function in the overall strategy space.*
That is,

$\frac{\partial^2 \pi_i}{\partial f_1 \partial f_2} < 0$ *for* $i \in \{1,\ 2\}$.

By changing the sign of one player's strategy space, we can trivially convert the game into a supermodular game. That is,
$\frac{\partial^2 \pi_i}{\partial f_1 \partial f_2} > 0$ *for* $i \in \{1,\ 2\}$.

[1]For the s-curve model, this holds for all cases where $\alpha < 2.4456$, a high value with respect to empirically estimated ranges, see [5]. See [9] for a discussion of parameter ranges for the schedule delay model.

Table 1: Approximated Concave Payoff Coefficient Estimates for Varied Seats per Flight

S	γ_0	γ_1	γ_2	γ_3	γ_4	γ_5	R^2
1000	122200	18135	-17856	-494	686	-533	0.96
250	122250	18130	-17861	-494	686	-533	0.96
225	122400	18115	-17876	-493	687	-532	0.96
200	122640	18095	-17901	-493	687	-532	0.96
175	123470	18030	-17989	-492	690	-529	0.96
150	125340	17925	-18214	-491	696	-523	0.96
125	129430	17885	-18838	-496	716	-514	0.95
100	136710	18277	-20301	-513	773	-518	0.94
75	142620	20224	-22355	-567	865	-578	0.93
50	104880	27814	-18929	-744	773	-846	0.93

Refer to [9] (working paper, link in reference) for proofs of these propositions. These results are significant because they demonstrate that subgame-perfect pure strategy Nash equilibrium is a credible and tractable solution concept for our two-stage game. In particular, the existence and uniqueness results indicate the suitability of pure strategy Nash equilibrium as a solution concept for the second-stage game. Concave payoffs, a property not guaranteed for one-stage models [7], ensure that individual first-stage payoff maximization problems are efficiently solvable and that a first stage equilibrium exists [13], and the supermodularity property ensures that several iterative learning dynamics converge to this equilibrium [11]. These analytic results are obtained for the aforementioned simplified model. Next, we extend them for more general game settings by relaxing each of these assumptions. We consider 1-, 2- and 3-player games; presence of a no-fly alternative; finite seating capacities; and connecting passengers. Our approach is to solve the second-stage fare game computationally, generating equilibrium fare decisions and profits for every set of frequency decisions for integer daily frequency values ranging from 1 to 20. Then, we fit quadratic approximations to these profits as functions of the frequencies of all players. For example, for a 2-player non-stop market with a no-fly alternative and with finite seating capacity, we approximate the payoff of airline 1 as

$$\pi_1 \sim \gamma_0 + \gamma_1 f_1 + \gamma_2 f_2 + \gamma_3 f_1{}^2 + \gamma_4 f_2{}^2 + \gamma_5 f_1 f_2$$

where f_i is the frequency of airline i and γ's are the coefficients to be estimated. Varying the parameters α (or φ and r in the schedule delay model), β, N and seating capacities over large ranges found in literature and practice, we find an excellent fit ($R^2 > 0.9$, ranging from 0.91 to 0.9998 in plausible parameter ranges) in all cases. The signs of all estimated coefficients are consistent with submodularity and concavity properties (e.g., $\gamma_3 < 0$ and $\gamma_5 < 0$ in the example above). For example, **Table 1** shows estimated payoff function coefficients for the approximated payoff function, with varied seats per flight S, $\alpha = 1.29$, $\beta = -.0045$ and N=0.5. In addition, estimated coefficients are consistent with a guaranteed-first state equilibrium according to the diagonal strict concavity condition for the parameters ranges tested, given by [13], with a few exceptions in the high values of $\alpha (> 1.7)$.

Note that every two-player submodular game is also su-

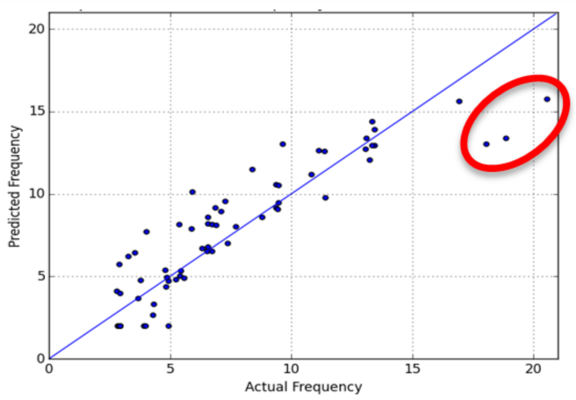

Figure 1: Predicted Versus Actual Frequency, without fixing high frequency segments

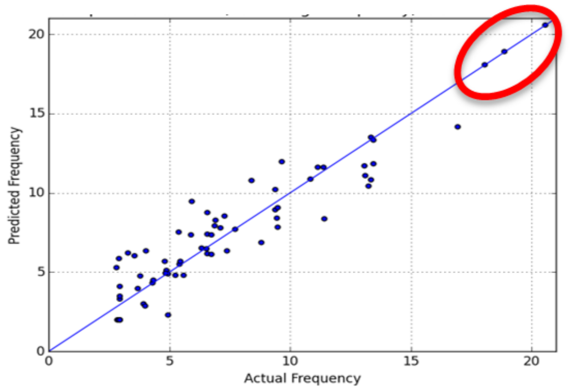

Figure 2: Predicted Versus Actual Frequency, fixing high frequency segments

permodular. Additionally, we confirm the following desirable property of quadratic-payoff games of used as approximations here (for its proof, by application of the generalized quasi-aggregative game framework in [10], refer to [9]).

PROPOSITION 4. *For an N-player game with concave quadratic submodular payoffs, the* myopic best response *heuristic, where each player optimizes its payoff against fixed opponent strategies iteratively, converges to a pure strategy Nash Equilibrium.*

Robustness of these concavity, submodularity and best response convergence properties across a wide range of parameter values and competitive scenarios suggests that the general model is highly tractable and efficiently solvable using the myopic best response heuristic.

3. AIRLINE NETWORK CASE STUDY

To test the tractability and predictive validity of our model in practice, we apply it to a 4-airline, 11-airport network in the Western U.S. The network includes 68 unique carrier segment pairs. We use route, flight cost, daily frequency, market demand and fleet data from the Bureau of Transportation Statistics database ([1], [2], [3]) from the first quarter of 2007. The 4 carriers playing the game are taken to be the 4 major carriers in the network - Southwest Airlines, United Airlines, US Airways, and Alaska Airlines. The frequencies of small regional carriers were held constant. We use the concave quadratic submodular payoff functions of airline frequency fitted to second stage payoffs in the above section and additionally enforce the aircraft availability constraints. Players thus allocate flight frequencies across their respective networks by solving a constrained quadratic program each best response iteration, continuing until convergence. The aforementioned *myopic best response* heuristic is found to converge to equilibrium in 6-7 iterations in <1 second of computational time. To calibrate the quadratic payoff coefficients of this model, we minimize the Mean Absolute Percentage Error ($MAPE$) of the predicted frequencies against actual frequencies:

$$MAPE = \frac{\sum_{cm \in CM} \left| \hat{f}_{\text{cm}} - f_{\text{cm}} \right|}{\sum_{cm \in CM} f_{\text{cm}}}$$

Here CM is the set of carrier-segments, and $\hat{f}_{\text{cm}}$ and f_{cm} are the predicted and observed frequencies respectively. We group carrier-segments into four categories (three-player markets, two-player hub-to-hub markets, other two-player markets, and one-player markets), with the carrier-segments within each group assigned a quadratic payoff function with identical coefficient values, and calibrate the resulting 11 payoff function coefficients using a stochastic gradient approximation algorithm from [14]), initializing with coefficients estimated for the following parameter values: no seating restrictions, $\alpha = 1.29$, $\beta = -.0045$ and N=0.5. The game is run repeatedly until convergence of coefficient estimates, with approximated $MAPE$ guiding coefficient updates at each step.

Figure 1 compares actual frequencies (x-axis) and predicted frequencies (y-axis) in the left panel after payoff coefficient calibration. The 45 degree blue line represents perfect predictions. Most predictions are found to be near this line. An overall MAPE of 18.4% is achieved, corresponding to 49% of absolute errors <1, and 78% <2. Notable outliers are the three highest frequency segments (circled in Figure 1), which are all flown by Southwest Airlines between its focus cities (hubs). Fixing these frequencies, rerunning the model, and omitting these frequencies from $MAPE$ calculations results in an $MAPE$ of 16.7% (Figure 2), suggesting that, except for the under-predictions of the highest frequencies, the model empirically performs well for most carrier-segments. We subsequently use this validated model to perform various scenario analyses. An expanded model for a network of 34 airports across the United States (OEP Airports, minus HNL), 9 airlines, and 545 carrier-segments, similarly converges to an equilibrium in 6-7 iterations in <1 second.

We also examine the out-of-sample performance of the model: in other words, we calibrate payoff coefficients on a past quarter, and make frequency predictions in a subsequent quarter, using demand, cost, market, and player attributes from that quarter as inputs to the model. For example, training our 11 payoff coefficients on data from Q1 of 2007 and testing on Q4 of 2007, we achieve an $MAPE$ of 20.6% (47% of absolute errors <1, and 72% <2). We can adjust our predictions for individual carrier-segments by the direction and magnitude of error for those carrier-segments in the training data (2007 Q1), if they exist. Using this procedure reduces our $MAPE$ to 11.2% (72% of absolute errors

Table 2: Model Performance in PDX-SFO Market, Q4 2007, Calibrated on Q1 2007 data

Carrier	True Frequency	Predicted Frequency
AS	3.02	3.52
UA	6.11	7.28

Table 3: Model Performance Q4 of 2007, Airport-Level Flight Count Predictions

Airport	Observed Q1	Observed Q4	Predicted Q4
LAX	137	128	132
SJC	60	64	61
LAS	154	167	168
SAN	101	108	110
SMF	70	71	70
SEA	98	105	104
PDX	33	41	43
SFO	67	92	99
ONT	61	61	62
PHX	159	159	153
OAK	88	88	85

<1, and 92% <2). Examining a market present only in our testing data, and not in our calibration data (PDX-SFO), we find that our model is still able to approximate individual observed frequencies (**Table 2**).

Our model also is able to approximate airline behavior at higher levels of aggregation. For example, we can examine the total number of flights associated with a certain airport, a quantity of interest for decision-makers concerned with airport congestion. **Table 3** shows observed flight counts at the airports in the network under study in Q1 and Q4 of 2007, and the number of flights predicted by our model (calibrated on Q1) for Q4. At airports where a significant increase in traffic levels was observed (e.g. LAS), our model was able to predict this increase.

4. CONCLUSIONS

This study investigates a two-stage frequency-fare game model which is behaviorally consistent with the actual airline decision process. For simple cases, we are able to prove various attractive analytical properties of this model indicating well-behaved and tractable games, with unique equilibria and favorable convergence properties. Using payoff function approximations, these properties are shown to extend to more realistic game settings. When applied to a real-world competitive network, the model converges quickly and hence is easy to calibrate. The model's predictions closely approximate actual frequency values, suggesting that refinements of the model could be pursued for use in scenario analysis, forecasting, planning, and policy making.

5. REFERENCES

[1] Air carrier financial reports (form 41). Office of the Assistant Secretary for Research and technology, Bureau of Transportation Statistics, United States Department of Transportation, Washington D.C. `http://www.transtats.bts.gov/tables.asp?DB_ID=135`. Accessed 2015-08-15.

[2] Air carrier statistics database (t100). Office of the Assistant Secretary for Research and technology, Bureau of Transportation Statistics, United States Department of Transportation, Washington D.C. `http://www.transtats.bts.gov/DatabaseInfo.asp?DB_ID=111`. Accessed 2015-08-15.

[3] Airline origin destination survey (db1b). Office of the Assistant Secretary for Research and technology, Bureau of Transportation Statistics, United States Department of Transportation, Washington D.C. `http://www.transtats.bts.gov/DatabaseInfo.asp?DB_ID=125&DB_Name=Airline%20Origin%20and%20Destination%20Survey%20`. Accessed 2015-08-15.

[4] N. Adler, E. Pels, and C. Nash. High-speed rail and air transport competition: game engineering as tool for cost-benefit analysis. *Transportation Research Part B: Methodological*, 4(10):812–833, 2010.

[5] P. Belobaba. *The Global Airline Industry*, chapter Overview of Airline Economics, Markets, and Demand. Wiley, 2 edition, 2016.

[6] J. Brueckner and R. Flores-Fillol. Airline schedule competition. *Review of Industrial Organization*, 30(3):161–177, 2007.

[7] 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.

[8] M. Hansen and Y. Liu. Airline competition and market frequency: a comparison of the s-curve and schedule delay models. *Transportation Research Part B: Methodological*, 78:301–317, 2015.

[9] R. Harder and V. Vaze. Two stage game theoretic modelling of airline frequency and fare competition. *Working Paper, `http://engineering.dartmouth.edu/mdsg/two_stage_game_frequency_fare.pdf`*, 2016.

[10] M. Jensen. Aggregative games and best-reply potentials. *Economic Theory*, 43(1):45–66, 2010.

[11] P. Milgrom and J. Roberts. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58(6):1255–1277, 1990.

[12] T. Morisset and A. Odoni. Capacity, delay and schedule reliability. *Transportation Research Record: Journal of the Transportation Research Board*, 2214:85–93, 2011.

[13] J. Rosen. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica*, 33(3):520–534, 1965.

[14] J. Spall. Overview of the simultaneous perturbation method for efficient optimization. *John Hopkins APL Technical Digest*, 19(4):482–492, 1998.

[15] V. Vaze and C. Barnhart. Modeling airline frequency competition for airport congestion mitigation. *Transportation Science*, 46(4):512–535, 2012.