

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

# Managing Air Traffic Disruptions Through Strategic Prioritization

Douglas Fearing

Technology and Operations Management, Harvard Business School, Boston, MA 02163, dfearing@hbs.edu

Ian A. Kash

Microsoft Research, Cambridge, UK, iankash@microsoft.com

The impacts of schedule disruptions in the U.S. air transportation system are substantial, with a recent study estimating the costs of congestion and delays at over \$30 billion for domestic operations in 2007. On the day of operations, if demand is expected to significantly outstrip available capacity (e.g., due to a severe storm), the Federal Aviation Administration implements air traffic flow management initiatives to safely resolve these imbalances. The most common measure enacted is a ground delay program, in which arrival slots into a congested airport are rationed to meet the projected capacity constraints. The assigned arrival slot determines the delay for each flight, which is realized as ground holding at the origin airport. The current approach for allocating arrival slots, ration by schedule, treats impacted flights equivalently regardless of the aircraft size, passenger load, mix of connecting passengers, etc. We extend this approach to develop a prioritized rationing scheme, *ration by prioritized schedule*, and show that significant benefits can be achieved through a prioritized allocation, even in the face of airline recovery responses. Subsequently, we develop a *strategic prioritization* game – a non-monetary, market-based scheme for allocating flight priorities that allows airlines to trade-off priorities across airports. In addition to demonstrating nice equilibrium properties, we show that our bidding and allocation scheme is capable of achieving some of the demand-management benefits of congestion pricing, which has been widely studied in the literature but has met with significant resistance in practice.

*Key words:* air traffic flow management, strategic prioritization, equilibrium analysis

*History:* This paper was first submitted on August 12, 2012.

---

## 1. Introduction

In the U.S. air transportation system, congestion and resulting delays place a tremendous financial burden on airlines, passengers, and the U.S. economy as a whole. The total cost of U.S. domestic air traffic congestion has been estimated as \$31.2 billion for the year 2007 (Ball et al. 2010). After

2007, congestion improved due to decreased customer demand (fewer passengers) and increasing load factors (fewer empty seats), but recent trends suggest we will soon reach and then exceed the levels of delay experienced during the 2007 peak.

One obvious approach for addressing the imbalance between capacity and demand is to increase the capacity of key system resources – the airports and the air sectors. For example, new runways could be added at the most congested airports, and en route separation requirements could be reduced through improved location technologies. Unfortunately, each of these approaches presents substantial challenges. It is difficult to build new runways where they are needed most because the most congested airports are typically in space-constrained, urban environments. Moreover, even when space can be found, new runways face substantial opposition due to political and environmental concerns. Improving separation requirements would yield only a modest increase in air sector capacity and would require upgrading the planes that use that sector.

Outside of capacity increases, there are, broadly, two approaches to addressing congestion. The first is to manage existing capacity more effectively, while the second is to incentivize airlines to schedule fewer flights. In our work, we show how to accomplish both of these objectives simultaneously through the *strategic prioritization* of flights. In strategic prioritization, we propose a scheme that forces airlines to make flight priority trade-offs at the time flights are scheduled (i.e., well in advance of operations). Intuitively, when there is a disruption and capacity needs to be rationed, the specified priorities allow the regulator to allocate capacity more effectively (i.e., to the airlines that value it most). Additionally, making these strategic trade-offs causes more of the congestion-related flight scheduling costs to be internalized by each airline, thus reducing over-scheduling.

The mechanism we use for incorporating flight priorities is air traffic flow management (ATFM). Air traffic flow management initiatives are implemented on the day of operations when there are expected to be significant imbalances between capacity and demand (e.g., due to a severe weather event). Ground holding programs are the most commonly applied ATFM initiatives. In a ground holding program, arrival slots for an airport or air sector are allocated to carriers, with expected arrival delays converted into ground holding at the airport of departure. These slots are allocated according to the posted flight schedule (i.e., first-scheduled, first-served) through a process known as Ration by Schedule (RBS). Note that this rationing treats all flights equivalently regardless of size or importance to the airline. Nonetheless, RBS has become the accepted view of fairness within the industry, and in the case of a single ground holding program, it has been shown to have nice theoretical properties relative to both fairness and efficiency (Vossen and Ball 2006).

A number of rationing schemes have been proposed as alternatives to RBS (e.g. Ball et al. (2010), Manley (2008)). Rather than proposing a wholly new rationing scheme, our approach slightly modifies RBS by essentially treating flights with priority as if they had been scheduled earlier

than their actual time. Since airlines are able to assign priorities in advance of operations, this effectively allows airlines to trade congestion costs across different airports and different days. Thus, our work is related to market-based mechanisms for ATFM, such as slot exchanges (Vossen and Ball 2006), auction-based slot allocation (Vazirani 2011), and day-of-flight waivers (Hoffman et al. 2011). Unlike slot exchanges and auction-based slot allocation, which are unlikely to be practically feasible due to the complexity introduced for airline recovery operations, our approach requires complex decisions in advance of operations but does not impact airline recovery approaches on the day of operations. Independent of our work, day-of-flight waivers have recently shown promise in a human-in-the-loop simulation setting. They represent another approach to prioritized slot allocation by allowing airlines to exempt individual flights. The key differences versus our approach are that strategic prioritization requires priorities to be set in advance, allows a spectrum of flight priorities rather than a binary exempt / non-exempt status, and is backed by an analysis of strategic behavior.

Since airlines have a limited budget of priority to assign, if they schedule fewer flights they can assign more priority to the remaining flights. Thus, airlines that choose to schedule more flights are forced to internalize some of the costs this imposes on other airlines' operations. Because of this, our approach realizes some of the benefits of market-drive approaches to congestion reduction such as slot auctions (e.g., Ball et al. (2006), Harsha (2008)) and congestion pricing (e.g., Brueckner (2002), Brueckner (2005), Morrison and Winston (2007)) However, unlike these market-based slot allocation approaches, our scheme is non-monetary, addressing airline concerns about the imposition of additional monetary costs on their operations. Opposition to monetary costs has been a significant barrier to previous attempts to reduce congestion through mechanisms such as a now-canceled auction of landing rights at New York-area airports (United States Department of Transportation May 13, 2009).

In the remainder of the paper we

- introduce the *Ration by Prioritized Schedule* (RBPS) algorithm for air traffic flow management (Section 2);
- provide simulation results derived from historical data showing that RBPS allows a more efficient allocation of congestion costs (Section 3);
- develop a game-theoretic model and equilibrium analysis of a system in which airlines strategically allocate these priorities (Section 4 and 5); and
- provide simulation results showing that this causes airlines to internalize and reduce system-wide congestion costs (Section 6).

Our first set of simulations takes historical data about flight schedules and ground delay programs and uses a detailed optimization model of airline recovery operations to estimate the effects of

RBPS on airline costs and passenger delays. The results suggest that even simple heuristics for assigning priority that could be implemented by the FAA today, for example assigning priority proportional to the number of seats on a flight, can lead to airline cost reductions of 4.0% and passenger delay reductions of 4.4%. Allowing airlines flexibility in allocating flight priorities leads to even larger improvements.

Giving airlines more flexibility in assigning priorities also introduces strategic considerations. Thus, we design a method of allocating priorities that guarantees the existence of pure strategy equilibria and avoids the need for airlines to randomize their priority allocations or constantly adapt to each others decisions. Furthermore, recognizing that the existence of a priority budget may have an effect on airline decisions about the number of flights to schedule, we analyze the subgame perfect equilibria of a two stage game where in the first stage airlines determine how many flights to schedule and in the second stage they bid for priorities. While there may be multiple pure strategy equilibria, we show that the total number of flights scheduled in all equilibria (weakly) decreases as the amount of prioritization provided increases. Thus, in addition to allocating delays more effectively, prioritization can be used to incentivize airlines to decrease the total amount of congestion in the system.

Our theoretical results are guaranteed to hold only for priority allocations of a limited size, but suggest that their conclusions may still hold in practice with larger allocations. Our second set of simulations provides evidence that this is the case, as well as helping quantify the extent to which this approach can close the gap between the number of flights scheduled under RBS and the smaller, socially optimal number of flights.

## 2. Rationing and Prioritization

The approach of allocating ATFM slots based on the scheduled flight order is widely accepted as fair within the industry. This Ration by Schedule (RBS) approach is utilized in both the U.S. and Europe. As an example of how this procedure works, the second column of Table 1 gives the scheduled arrival times<sup>1</sup> of three hypothetical flights. If bad weather causes an ATFM program to be put into force where flights are limited to one arrival per 20 minute slot, the resulting RBS schedule is given by the fourth column.

One critique of the RBS approach is that, outside of some distance-based exemptions, there is no differentiation between flights. Thus, a Boeing 777 with approximately 300 seats is slotted in the same manner as an Embraer 190 with just 100. In addition to causing excessive passenger delays on the day of operations, this approach creates a mild incentive for airlines to schedule more flights on smaller aircraft in order to garner a larger share of ATFM slots, because airlines are permitted

Flight	Scheduled	Prioritized	RBS	RBPS
A1	07:00	06:49	07:00	07:00
B1	07:15	07:10	07:20	07:40
A2	07:20	07:09	07:40	07:20

**Table 1** Slot rationing comparison for two airlines with slot sizes of 20 minutes starting at 07:00.

to reallocate these slots among their flights. This runs directly counter to the incentives regulators would like to provide in order to reduce over-scheduling and congestion.

To address these concerns, we describe the Ration by Prioritized Schedule (RBPS) allocation scheme, a simple extension of the RBS approach used in practice. In our RBPS allocation, each flight  $f$  is assigned a priority  $p_f$ , and slots are rationed according to the priority-adjusted schedule. That is, assuming flight  $f$  is scheduled to arrive at  $\alpha_f$ , under RBPS slots would instead be allocated in increasing order based on  $\alpha_f - p_f$ . Thus, our approach is equivalent to applying the RBS allocation based on the priority-adjusted schedule.

The remaining columns of Table 1 show how RBPS might differ from RBS in our example. For the RBPS allocation, airline A’s flights are assigned a priority of 11 minutes each, and airline B’s flight is assigned a priority of 5 minutes. Due to the difference in flight priority, flight A2 is able to depart on-time under the RBPS allocation, pushing all of the delay to flight B1. It is worth noting that the delays assigned to flights B1 and A2 are discontinuous with respect to the priorities. Reducing the priority of flight A2 by more than minute causes the RBPS allocation to revert to the RBS ordering, shifting 20 minutes of delay from airline B to airline A, whereas any reduction less than a minute leads to no change. This issue will return when we discuss strategic prioritization in Section 4.

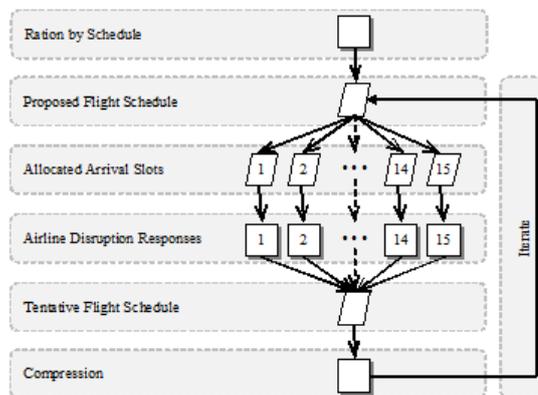
RBPS is not the only way to incorporate priorities into the RBS procedure. Day-of-flight waivers (Hoffman et al. 2011) provide a binary approach which can be thought of as RBPS where most flights are given no priority but a few are given infinite priority, subject to the constraint that they cannot arrive earlier than their scheduled arrival time. An alternative prioritization approach would be to treat flight priorities as multiplicative factors in a lexicographic delay cost minimization objective (RBPS can be interpreted as incorporating additive factors in this same optimization). The practical concern with a multiplicative approach is that the resulting deviation from the scheduled order is unbounded; in the presence of severe delays, the scaled delay cost objective would determine the allocation order based solely on the priorities, ignoring the original schedule except to break ties. In comparison, our approach guarantees that a flight  $f$  will not jump ahead of any flights scheduled to arrive more than  $p_f$  minutes in advance of it.

Thus far, we have said nothing about how priorities should be allocated to individual flights. In the remainder of the paper, we show ways that priority can be allocated to achieve two different set

of goals. The next section examines policies for assigning priorities that can better allocate delays among flights to reduce both overall airline costs and overall passenger delays (although the total amount of congestion in the system is unchanged). In the policies we compare, the regulator plays a substantial role in determining the points to allocate among either airlines or flights, which assumes some knowledge of underlying costs and potential benefits. After that, we examine a strategic prioritization model where airlines determine flight priorities through a non-monetary, bidding process, and show that the resulting tradeoffs cause airlines to schedule fewer flights, reducing the overall congestion in the system.

### 3. Simulations of Prioritization

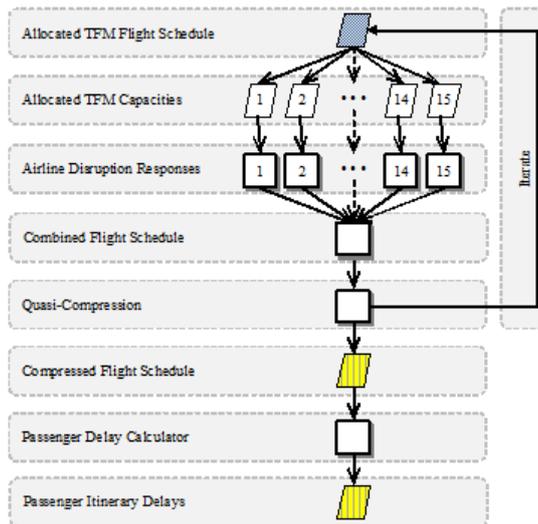
In this section we examine a number of policies for assigning priorities and show that they lead to lower costs for both airlines and passengers. One challenge with evaluating any prioritized allocation scheme is the complex dynamics that occur after slots are rationed. In the U.S., ground delay programs (GDPs) are managed according to a procedure known as Collaborative Decision Making (CDM). Within CDM, airlines are free to adjust flight schedules within their rationed slots by canceling or reordering their flights. This ensures that the resulting schedule is feasible, although each flight cancellation creates a gap in the schedule. To maximize utilization of congested resources, the FAA performs Compression to collapse the schedule through a sequence of 1-to-1 slot exchanges. The overall CDM process is depicted in Figure 1 and described in detail in (Chang et al. 2001).



**Figure 1** Depiction of Collaborative Decision-Making (CDM) procedure.

Through the CDM process, because slots can be treated as unspecified arrival capacity, each airline is afforded significant flexibility to prioritize its flights within a single GDP. Within this environment, evaluating a prioritized allocation without considering the airline response is problematic, because airlines incorporate priorities internally when making recovery decisions. Thus,

for the purposes of our experiments, we extend a sequential evaluation framework developed in Fearing and Barnhart (2011) to evaluate the benefits and key considerations associated with our RBPS allocation. The framework attempts to mimic the iterative CDM process by performing optimization-based approximations of the decisions made by the airlines and the FAA. To model airline disruption responses, the framework utilizes an integer optimization model that balances operational and passenger considerations to make determinations on which flights to cancel, swap, or delay. Subsequently, a quasi-compression is performed to fill any gaps created in the schedule. As with the CDM process, our method iterates until convergence is reached (i.e., when airlines make no additional changes to the compressed schedule). Each allocation policy is evaluated by comparing recovery objective values and estimating passenger delays associated with the resulting schedule. Passengers delays due to missed connections or flight cancellations are estimated by (greedily) re-accommodating each of the disrupted passengers to his or her destination on a future flight (or flights) with available seats. The sequential evaluation procedure is depicted in Figure 2.



**Figure 2** Depiction of sequential evaluation procedure with inputs in blue and outputs in yellow.

For our experiments, we consider the same 10 historical scenarios that are developed in Fearing and Barnhart (2011). Each of these scenarios is created by overlaying historical GDPs on a single day of clear weather operations. The operational data include only the executed flight schedules, so the clear weather results are used to provide a reasonable estimate of the pre-disruption schedule. The data for our scenarios was provided by Metron Aviation, a company that develops software used by the FAA to manage ATFM programs. Table 2 summarizes key information for each scenario including the date of the historical disruption, the airports where GDPs have been initiated, and the number of flights impacted by GDPs.

Date	Airports	Flights
05/02/2007	(2) LGA SFO	570
05/09/2007	(2) IAD JFK	480
06/19/2007	(8) ATL DCA EWR IAD JFK LGA SFO PHL	2522
06/27/2007	(4) CYYZ JFK LGA ORD MDW	1896
06/28/2007	(5) EWR IAD JFK LGA SFO	1398
07/05/2007	(2) CYYZ EWR	348
07/16/2007	(2) LGA SFO	570
07/18/2007	(5) EWR JFK LGA PHL SFO	1320
07/27/2007	(4) EWR LGO ORD SFO	1726
09/27/2007	(6) ATL CYYZ EWR JFK LGA PHL	1836

**Table 2** Summary of key information for historical ATFM scenarios.

Scope Criteria	None None		All Seats		Airline Delay x Seats		Airline x Airport Delay x Seats	
	Obj (1000s)	Px Dly (days)	Obj (1000s)	Px Dly (days)	Obj (1000s)	Px Dly (days)	Obj (1000s)	Px Dly (days)
Date								
05/02	199	411	195	408	174	359	191	409
05/09	73	162	59	136	54	128	60	138
06/19	2245	3505	2100	3274	2010	3236	2053	3152
06/27	1723	2796	1680	2716	1626	2658	1674	2691
06/28	1719	3015	1685	2908	1457	2441	1658	2893
07/05	10	72	9	69	9	68	9	70
07/16	115	256	109	221	83	200	97	218
07/18	1114	1945	1093	1925	1111	2079	1070	1887
07/27	306	514	278	469	260	422	276	467
09/27	640	1030	608	969	616	1010	606	973
<b>Total</b>	8145	13705	7815	13095	7399	12602	7695	12898

**Table 3** Results from simulation of various prioritization policies. For each policy and scenario, we calculate the sum of the recovery objective values (Obj) in thousands and the total passenger delays (Px Dly) in days.

In our results, we compare four allocation policies, each specifying a different approach for allocating priority minutes to airlines and assigning these minutes to individual flights. Across all of the policies we consider, we hold the total number of prioritization minutes constant, equal to fifteen minutes per flight. In Table 3, we report on both the sum of the recovery objective values and the total passenger delays for each scenario and policy. The two reported statistics represent the airline recovery costs and passenger benefits, respectively. It is worth noting that the underlying airline response model maintains aircraft flow balance, ensuring that flights are canceled in cycles to satisfy positioning requirements for future operations. This ensures that canceled flights do not impose hidden costs by requiring planes to be relocated before normal operations can resume.

In our first comparison, we consider two centrally-planned allocation policies. The first is unprioritized, applying the traditional RBS procedure (None - None in Table 3). The second applies RBPS, assigning each flight a priority proportional to the number of seats on the plane (All -

Seats). In aggregate, prioritizing the flights proportional to the number of seats results in a 4.4% reduction in passenger delays and a 4.0% reduction in airline recovery costs.

In our second comparison, we compare a centrally-planned prioritization policy (All - Seats) to one where the assignment of priorities to flights is controlled by each airline (Airline - Delay x Seats). To determine the airline-specific flight priorities, we apply a two-stage approach. First, the central planner allocates each airline a prioritization pool equal to the aggregate number of minutes it would be allocated under the centrally-planned approach. Second, the airline adjusts these priorities, assigning each flight a priority proportional to a proxy for internal costs. The proxy we use for these costs is the average delay allocated across all scenarios according to RBS (a measure of congestion) times the number of seats on the plane (a coarse measure of the operating costs). In aggregate, even based on this relatively coarse proxy for internal costs, we see a substantial benefit associated with airline-specific assignment of flight priorities. In particular, passenger delays have been reduced by an additional 3.8%, and airline recovery costs have fallen by 5.3%.

In our third comparison, we consider the scope of the airline's priority assignment decisions. In the first policy, the prioritization decisions are made across all flights (Airline - Delay x Seats), whereas in the second, the prioritization decisions are made at each individual airport (Airline x Airport - Delay x Seats). Within each of these scopes, the airline assigns priorities to its flights proportional to the cost proxy defined above. In aggregate, we see that preventing airlines from reallocating priority between airports leads to an increase of passenger delays of 2.4% and an increase in airline recovery costs of 4.0%.

In summary, our simulation results suggest that:

- prioritized slot allocation could significantly reduce passenger delays and airline recovery costs;
- flight priority decisions are best made by the airlines due to their visibility into internal costs;
- and
- there are substantial benefits to allowing airlines to trade-off priorities across airports, a capability that does not exist in the current system.

These insights motivate the development of our strategic prioritization model in the following section.

#### **4. Model of Strategic Prioritization**

Thus far, we have explored the Ration by Prioritized Schedule (RBPS) under the assumption of non-strategic behavior by airlines. Since flight schedules were kept fixed, the total delay experienced by flights (before accounting for cancellations), was also fixed. However, we saw that RBPS could change which flights experienced the delays, to the benefit of both airlines and passengers. For the remainder of the paper, we focus on a the ability of RBPS to affect the total amount of delay

by influencing the number of flights scheduled. To do so, we explore a greatly simplified model that allows us to focus on strategic behavior, while ignoring the differences in importance between different flights. Using this model, we show that prioritization decisions force airlines to internalize more of the costs of the flights they choose to schedule, resulting in an overall lower number of flights and greater social welfare.

Before formally describing our model, we present an informal overview. Airlines participate in a two stage game. In the first stage, airlines determine how many flights to schedule at each airport. They receive positive utility for each flight they schedule but negative utility for the expected delays these flights experience due to congestion. These congestion costs are determined in the second stage, in which airlines acquire priorities at each airport leading to a RBPS allocation. In our design of a mechanism for determining priorities, we desire an approach that 1. guarantees the existence of a pure strategy equilibrium; 2. requires airlines to make prioritization trade-offs across airports; and 3. allows the regulator to determine the influence the mechanism has on resulting schedules. The existence of a pure strategy equilibrium is beneficial, because we do not believe mixed strategies are realistic in this environment. In order to reduce over-scheduling, our mechanism should force airlines to internalize some of the congestion-related externalities. This can be accomplished by forcing airlines to make prioritization trade-offs across airports, implicitly pricing the congestion costs at each one. Last, for practical viability it is helpful to consider a mechanism that allows regulators to vary the relative influence, both across airports and over time, in order to support a staged implementation.

With these objectives in mind, we propose the following mechanism for *strategic prioritization*. First, for each airport, the regulator specifies a pool of prioritization minutes which will be split among all flights scheduled to arrive at that airport. Having a separate prioritization pool for each airport provides control over the extent to which priorities can affect the corresponding operations. Each airline has a pool of points to spend bidding for priority. Airlines acquire a fraction of the priority pool at an airport that is proportional to the number of points they bid there. Such proportional allocation schemes have been studied in a number of contexts (see, e.g., Kelly et al. (1998)) and often result in stable, unique equilibria. However, our setting introduces two analytical complications. First, rather than spending money, airlines are spending a fixed pool of points across multiple such auctions, which requires a more complex equilibrium characterization. Second, for a fixed scenario, the value of assigning priority to a flight is a step function, with each jump occurring when the flight acquires enough priority to move ahead of another flight. Such discontinuities lead to problematic behavior in many settings, but because we consider prioritization on a strategic timescale, where decisions are made well in advance of operations, we can take advantage of randomness which results in a continuous expected value function. We adopt a model where the

actual scheduling on a day is randomly determined, so at the time of bidding for priority, airlines do not know which flights will be scheduled ahead of others. This yields an analytically tractable model.

We now present our formal model of the game. There are two airlines and at least one airport. In the first stage of the game, each airline  $a$  simultaneously determines the number  $N_{ra} \geq 1$  of flights to schedule at each airport  $r$ . In the second stage, airlines bid for a share of  $P_r$ , the priority pool at airport  $r$ , using a proportional allocation mechanism. Note that we consider only the case where each airline schedules at least one flight at each airport, because the second stage is uninteresting if one airline schedules zero flights. The share of prioritization minutes each airline acquires is spread evenly among its flights, so if each airline  $a'$  bids  $b_{ra'}$ , then the priority assigned to flight  $i$  of airline  $a$  at airport  $r$  would be

$$p_{rai} = \frac{P_r}{N_{ra}} \frac{b_{ra}}{\sum_{a'} b_{ra'}} \quad (1)$$

Thus, under this framework, airlines are able to vary priorities across airports, although not across flights within a single airport. Airlines bid not with actual money, but with points from a pool of size  $B_a$ . Thus, their bids are subject to a budget constraint that  $\sum_r b_{ra} \leq B_a$ . In practice, this budget would likely be determined based on historical operations (e.g., flights, seats, or passengers carried). Additionally, we constrain each airline's bid  $b_{ra}$  to be at least some small, fixed  $\epsilon > 0$  to avoid technical complications introduced by permitting zero bids.

After airlines make their decisions, each flight  $i$  of airline  $a$  at airport  $r$ , gets a scheduled arrival time  $\alpha_{rai} \in [0, T_r]$  drawn uniformly at random. Subsequently, at each airport  $r$ , there is good weather with probability  $1 - \lambda_r$ , in which case all flights land on time. In the absence of over-scheduling, this can be considered equivalent to an airport operating in visual meteorological conditions (VMC). Otherwise, the delays due to bad weather take the form of a ground holding program, in which there is a fixed slot size  $z_r$  that determines how much time each arrival requires. That is, the first flight lands at time 0, the second at time  $z_r$ , the third at  $2z_r$ , etc. Based on the RBPS allocation, the actual arrival time of flight  $i$  of airline  $a$  at airport  $r$  is

$$\tilde{\alpha}_{rai} = z_r \sum_{(r,a',j) \neq (r,a,i)} \mathbb{I}(\alpha_{ra'j} - p_{ra'j} \leq \alpha_{rai} - p_{rai}) \quad (2)$$

Note that  $\tilde{\alpha}_{rai}$  could be less than  $\alpha_{rai}$ . In practice, this would likely be infeasible, but we assume that this is acceptable (and in fact beneficial) to avoid introducing additional analytical complexity due to the compression mechanism that would be used to resolve this issue in practice. Thus, the flight experiences a delay of  $d_{rai} = \tilde{\alpha}_{rai} - \alpha_{rai}$ . The final utility received by airlines is  $V$  for each flight they schedule minus the delay cost  $d_{rai}$  for each flight. Thus, the expected utility airline  $a$  receives for flight  $i$  at airport  $r$  is

$$u_{rai} = \mathbb{E}_{\alpha_r} [V - \lambda_r d_{rai}]. \quad (3)$$

$N_{ra} \geq 1$	number of flights of airline $a$ assigned to airport $r$
$-a$	the airline that is not airline $a$
$P_r$	bad-weather prioritization pool for airport $r$
$B_a$	budget constraint for airline $a$
$b_{ra} \geq \epsilon > 0$	bid by airline $a$ for priority at airport $r$ (required to be at least $\epsilon$ )
$p_{rai}$	priority of flight $i$ of airline $a$ at airport $r$
$\lambda_r$	bad-weather probability for airport $r$
$T_r$	duration of peak period at airport $r$ in minutes
$\alpha_{rai} \sim \mathcal{U}(0, T_r)$	scheduled arrival time for flight $i$ of airline $a$ at airport $r$ (uniformly distributed)
$\tilde{\alpha}_{rai}$	actual arrival time for flight $i$ of airline $a$ at airport $r$
$z_r$	bad-weather slot size at airport $r$
$d_{rai}$	delay for flight $i$ of airline $a$ at airport $r$
$V$	value of scheduling a flight
$u_{rai}$	expected utility due to flight $i$ of airline $a$ at airport $r$
$v_a$	expected utility of airline $a$ across all airports and flights

**Table 4** Summary of Notation

with the total expected utility being the sum across all flights  $v_a = \sum_r \sum_{i=1}^{N_{ra}} u_{rai}$ . Note that this formulation assumes that delay costs are linear. In general, we would expect them to be a convex, increasing function, a point we discuss further in Section 7.

Table 4 summarizes our notation.

## 5. Analysis of Strategic Prioritization

In this section, we analyze the equilibria of our model. Our main result is a comparative statics result, which shows that the pool size  $P_r$  is a natural tool to control the incentives of airlines: subject to an upper bound, increasing  $P_r$  (weakly) decreases the number of flights airlines schedule in equilibrium. This result has the caveat that for particular values of  $V$  some airlines may be indifferent between scheduling two different numbers of flights in equilibrium. However, such indifference disappears with an arbitrarily small change to  $V$ , so we say that it holds for "almost all  $V$ ". In standard game theoretic terminology the theorem holds *generically*. That is, it holds except at a subset of the parameter space of measure zero. The theorem focuses on pure strategy subgame perfect equilibria. As previously discussed, mixed strategies do not seem plausible in this setting. Subgame perfection is a condition that ensures that airline beliefs about what would occur if they scheduled different numbers of flights are based on the assumption that other airlines behave strategically. After stating the theorem, we give key lemmas used to prove it, and outline the arguments behind them, with the proofs deferred to the appendix.

**THEOREM 1.** *If  $P_r \leq \frac{T_r}{2}$  at all airports  $r$ , then for almost all  $V$ :*

1. *this game has a pure strategy subgame perfect Nash equilibrium;*

2. for all airports  $r$ , if there exists such an equilibrium where  $N_{r1}$  and  $N_{r2}$  flights are scheduled at  $r$  and another where  $N'_{r1}$  and  $N'_{r2}$  are, then  $N_{r1} + N_{r2} = N'_{r1} + N'_{r2}$ ; and
3. this equilibrium number of flights is nonincreasing in  $P_r$ .

Because this game has two stages, we begin our analysis with the second stage and show there exists a pure strategy equilibrium, and furthermore that all pure strategy equilibria have equivalent payoffs. That is, we show that there exists a unique equilibrium payoff in the second stage game. In the second stage, each airline chooses a bid  $b_{ra} \geq \epsilon$  at each airport. As  $d_{rai}$ , flight  $i$ 's delay, depends on the priorities of all flights at airport  $r$ ,  $u_{rai}$  is a function  $u_{rai}(\mathbf{b}_r)$  of the vector of bids at airport  $r$ . Our first lemma shows that there are diminishing marginal returns for bidding more at an airport as long as the size of the priority pool is sufficiently small.

LEMMA 1. For all airlines  $a$ , airports  $r$ , flights  $i$ , numbers of flights  $N_{ra}$  and  $N_{r-a}$ , and all bids of the other airline  $b_{r-a}$ ,  $u_{rai}(b'_{ra}, b_{r-a})$  is a continuous, strictly concave, and strictly increasing function of  $b'_{ra}$  if  $P_r \leq \frac{T_r}{2} \min(N_{ra}, N_{r-a})$ .

Most of the proof is straightforward, with the only difficulty being the proof of concavity. As we are using a proportional allocation mechanism (Kelly et al. 1998), it is immediate that there are diminishing marginal returns on the amount of priority gained, but additional analysis is needed to show that this translates into diminishing marginal returns on utility.

Continuity and concavity are sufficient for the existence of an equilibrium, but not for uniqueness. However, the special case we are analyzing is a two-player constant-sum game: the total delay cost is fixed and bidding determines how it will be split between the airlines. Thus, we can apply Sion's minimax theorem (Sion 1958) (a generalization of the standard von Neumann minimax theorem to games with infinite strategy sets) to show that equilibrium payoffs are unique. Furthermore, under the condition that each airline is allotted the same budget, we can construct an equilibrium to determine what these payoffs are. In this equilibrium, both players make the same bids (and thus each gets half the pool at each airport).

LEMMA 2. For all airports  $r$  and numbers of flights scheduled,  $N_{r1}$  and  $N_{r2}$ , if for all  $r$   $P_r \leq \frac{T_r}{2} \min(N_{r1}, N_{r2})$ , there exist unique equilibrium payoffs  $(v_1, v_2)$  in the second stage game. Furthermore, if each airline has the same budget ( $B_a = B_{-a}$ ) and the minimum bid  $\epsilon$  is sufficiently small, these payoffs are equivalent to those realized when both airlines choose the same strategy (when for all airports  $r$ ,  $b_{ra} = b_{r-a}$ ).

Because the second stage always has a unique equilibrium payoff, to find subgame perfect equilibria we can analyze the first stage as a normal form game whose payoffs are the unique equilibrium

payoffs of the resulting second stage game. This game is supermodular (has strategic complementarities) (Milgrom and Roberts 1990). In general, this means that when one player increases its strategy the best reply for the other player is nondecreasing. In a two player game, this condition also holds when the best reply is nonincreasing by reversing the ordering of strategies for one of the two players. In our context, this means that when one airline schedules more flights, the optimal number of the other airline to schedule (weakly) decreases.

LEMMA 3. *If for all  $r$ ,  $P_r \leq \sqrt{8/3}T_r$  and in all second stage games both airlines use the same strategy, then the first stage game is supermodular.*

Topkis (Topkis 1979) showed that, in games with strategic complementarities, Tarski's fixed point theorem (Tarski 1955) can be used to prove the existence of a pure strategy Nash equilibrium. This completes the proof of the first part of our main result as  $P_r \leq \frac{T_r}{2}$  satisfies the conditions of all lemmas. The remaining two parts follow through an analysis of airlines' marginal costs for adding flights.

LEMMA 4. *If for all  $r$ ,  $P_r < \sqrt{6}T_r$ , then for almost all  $V$ , all pure strategy subgame perfect equilibria where both airlines use the same strategy in the second stage and the resulting first stage game is supermodular have the same number of flights  $N_r = N_{r1} + N_{r2}$  at each airport  $r$ . Furthermore, if  $P_r \leq 2T_r$  for all  $r$ ,  $N_r$  is nonincreasing in  $P_r$ .*

Note that these lemmas could be stated more concisely if we simply took  $P_r \leq \frac{T_r}{2}$  in each of them. However, most of the bounds established in the proofs of the lemmas are sufficient but not necessary. Thus, we state them in these more general forms to remind the reader that they may well continue to hold for larger values of  $P_r$ . We examine this issue numerically in the next section.

## 6. Simulations of Strategic Prioritization

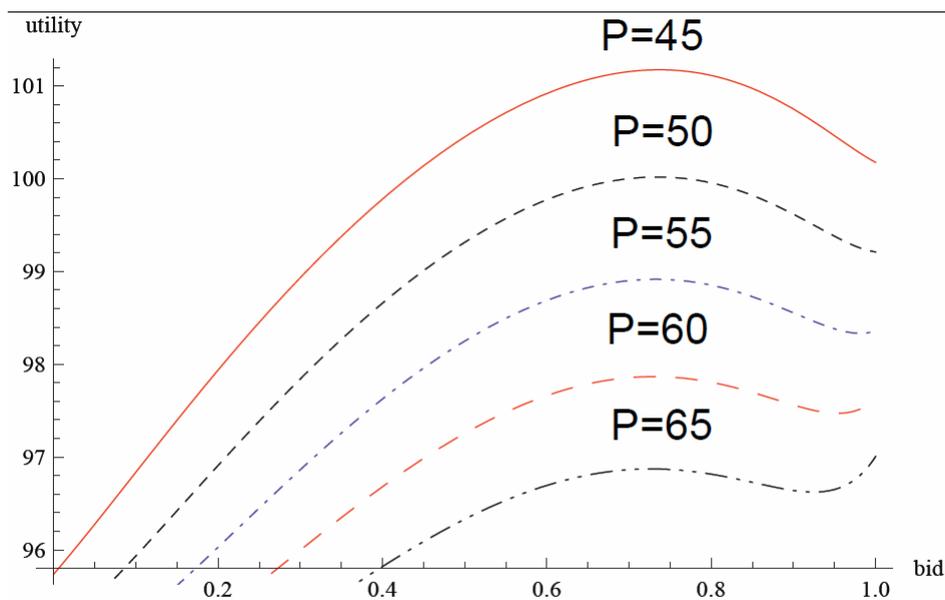
Our theoretical results established two key points. First, if  $P_r \leq T_r/2$ , a pure strategy subgame perfect equilibrium exists. Second, subject to this constraint, increasing  $P_r$  does not increase the equilibrium number of flights scheduled. In this section, we consider scenarios which show that, in practice, we would expect these two claims to hold for even larger values of  $P_r$ . Furthermore, our scenarios show that, beyond simply not increasing, the number of flights tends to decrease and approach the social optimum. The parameters for our scenarios are summarized in Table 5

Lemma 1 showed that utilities are concave in bids for  $P_r \leq T_r/2 \min(N_{ra}, N_{r-a})$ . Figure 3, which plots the utility of an airline for different bids and pool sizes<sup>2</sup>, shows that this is not necessarily true for larger values of  $P_r$ . However, our application of concavity in the proof of Lemma 2 is to guarantee that a bid that is locally optimal is also globally optimal. For this quasiconcavity is sufficient, which holds at least through  $P_r = 50$ . Although  $P_r = 60$  is not quasiconcave, the

Parameter	Value
Number of airports	2
$N_{ra}$	$\{1 \dots 10\}$
$P_r$	varies, but same for both airports
$B_a$	1
$\epsilon$	$10^{-9}$
$\lambda_1$	0.8
$\lambda_2$	0.4
$T_r$	60 for both airports
$z_r$	varies in $\{5, 10, 15, 20, 25, 30\}$ , but same for both airports
$V$	29

**Table 5** Summary of simulation parameters

maximum is still at the desired interior point. We have verified that, for all values of  $z_r$  used in our scenarios, this property holds through at least  $P_r = T_r$ , even in the worst case where one airline has only a single flight scheduled at each airport. However, as Figure 3 shows, this property breaks down for slightly larger values, and the characterization in Lemma 2 that airlines submit the same bids no longer applies.



**Figure 3** With larger  $P_r$ , utility is not concave.

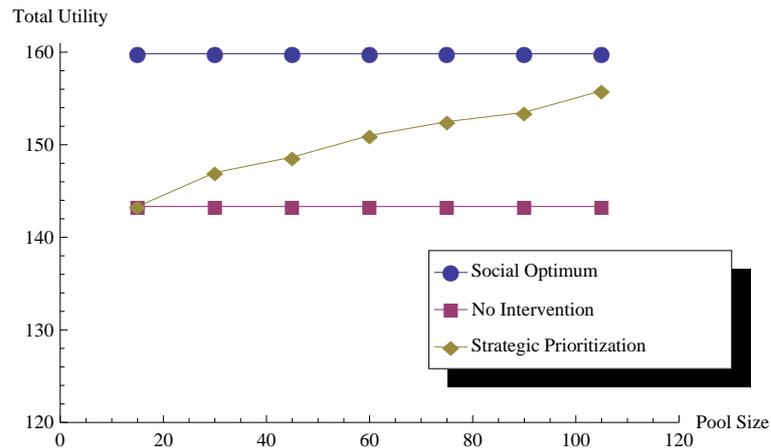
Our experimental results demonstrate, that in these scenarios, our game is supermodular and the number of flights scheduled is weakly decreasing in  $P_r$ , at least through  $P_r = T_r$ . However, note that even past this point concavity fails only when an airline schedules at most 2 flights at

some airport. If both airlines schedule a larger number of flights, then the results of Lemma 4 will continue to hold (at least up to its bound of  $\sqrt{6T_r}$ ). Even this bound is based on an assumption that the number of flights is small, so in general it is plausible that pool sizes that are linear in the equilibrium number of flights could be tolerated. The parameters we have chosen ensure that the equilibrium number of flights is small, which limits the computational costs of verifying that the bids from Lemma 2 are in equilibrium for all relevant second stage games. Thus, these parameters are not suitable for investigating larger pool sizes.

The results of our scenarios, aggregated across choices of  $z_r$ , are presented in Figure 4. The no intervention line is the result of priority pools of size 0, and roughly captures how the system works today. The social optimum specifies the best solution in terms of the total value across carriers (e.g., the best achievable through congestion pricing). The value of strategic prioritization varies between no intervention and the social optimum based on the size of the prioritization pool. To compute the no intervention and social optimum results for each scenario, we simply enumerate all 10,000 possible strategy pairs. For strategic prioritization, we exploit the fact that best-reply dynamics starting from an extreme point are guaranteed to converge to an equilibrium in supermodular games. Thus, we assume that all second-stage games have equilibria of the form given by Lemma 2 and make best replies until a candidate equilibrium is reached. At this point, we can verify that actual equilibria of the relevant second stage games support this candidate as a subgame perfect equilibrium. The last data point is for  $P_r = 105$ , because this procedure failed for some values of  $z_r$  with  $P_r = 120 = 2T_r$  and we have not tried other approaches (e.g., solving all second stage games and looking for equilibria in the resulting normal form game). The results suggest that strategic prioritization can capture a significant fraction of the benefits of other approaches such as congestion pricing, even it is not capable of achieving the socially optimal outcome.

## 7. Discussion

We have seen that the current approach for rationing ATFM capacity, Ration by Schedule, treats impacted flights equivalently regardless of the aircraft size, passenger load, mix of connecting passengers, etc., and that this means that flight delays are allocated inefficiently. Our simulations based on historical data show that our proposed Ration by Prioritized Schedule approach can achieve improvement even in the face of airline recovery responses. We have developed a strategic, non-monetary scheme for allocating flight priorities which allows airlines to trade-off priorities across airports. We have shown both theoretically and numerically that, in addition to having nice equilibrium properties, our bidding and allocation scheme is capable of achieving some of the benefits of congestion pricing, which has been widely studied in the literature but has met with



**Figure 4** Comparison of social welfare across three approaches with varying pool sizes.

significant resistance in practice. While these initial results are promising, there is still significant work to be done.

One avenue to explore results from taking the perspective of the regulator (i.e., the FAA in the U.S.). Our simulation results suggest that even simple heuristics, such as allocating priority proportional to the number of seats on an aircraft, can reduce congestion costs even after taking into account airline recovery procedures. Can such heuristics be theoretically justified or are there other rules the regulator could implement that can be? One natural candidate is the allocation derived from our equilibrium analysis: allocate each airline a pool of prioritization minutes at each airport (perhaps based on historical operations) to be divided evenly among its flights. This type of approach has the advantage that, as there is no decision making by airlines with regards to priority allocation, it would be simple to implement. On the other hand, any such mechanical rule is likely to be better for some airlines than others, which could create difficulty in reaching consensus on what procedure to adopt. Further, our historical simulations have shown that the change in slot allocation has benefits when airline recovery is taken into account, but, as the scenarios are based on historical flight schedules, they do not take into account how airlines would change their scheduling and operations as a result of a priority allocation rule. It is possible that airline scheduling responses would improve our results further, but additional analysis is needed.

Even if a strategic prioritization approach is used where airlines are the ones responsible for making priority allocation decisions, there are still interesting issues from the regulator’s perspective. For example, how should the pool size  $P_r$  be set at different airports? Should the size vary depending on the cause or severity of the ATFM program, or perhaps as a way of introducing randomness similar to schedule uncertainty in our simple model? As airlines vary substantially in size and affiliations (e.g., regional carriers and code shares), how should each airline be allocated a

points budget to bid for priority? In addition to these issues that are apparent even in our simple theoretical model, there are likely to be many practical market design issues. For example, how often and how close to the day of operations should airlines be allowed to change their bids? On one hand, the closer to day of operations airlines are allowed to make changes, the better the sense they are likely to have of their valuations. On the other, if bids can be changed too close to operations, there is likely to be little uncertainty so the discontinuous value of priority could lead to instability.

Turning to our theoretical analysis of airline behavior, our results made heavy use of the characterization of the second stage game as two player, constant sum, and symmetric. Many natural generalizations would violate these assumptions. If we specify different points budgets for each airline, we would need to characterize a different equilibrium bidding strategy. If there are more than two airlines, the game is still constant sum among them, but minimax theorems typically only apply in two player settings. If we allow for delay costs that are represented by a more general convex increasing function rather than linear or that are not the same for each airline, this results in a game that is not constant sum. Extending to these cases requires new analysis of whether there is still a unique equilibrium result. One promising approach is to find conditions where the utility functions satisfy a stronger condition than concavity known as diagonal strict concavity, which Rosen showed is a sufficient condition for the existence of a pure strategy unique equilibrium in this class of game (Rosen 1965). Extending beyond two airlines will also require a different analysis of the first stage game because we exploited the fact that we could reverse the natural ordering on one player's strategies to make the resulting game supermodular.

Finally, our strategic prioritization approach requires airlines to submit a single bid at each airport that is then divided evenly among each flight. In reality, some flights are more valuable than others, and it would be helpful to allow airlines to submit bids for each flight individually. The initial part of the analysis still goes through: there are diminishing marginal returns for bidding more on a single flight. However, the equilibrium analysis becomes more complex for both the second stage bidding game and the first stage scheduling game. Furthermore, in such a game, the ability for an airline to reallocate the slot gained by prioritizing one flight over another is likely of significant strategic importance.

## Appendix. Proofs from Section 5

### Proof of Lemma 1

We begin with several definitions that narrow the function we need to analyze. Let  $\tau_{ra}(\mathbf{b}_r)$  be the priority a flight of airline  $a$  has over a flight of airline  $-a$ . This has the value

$$\tau_{ra}(\mathbf{b}_r) = \frac{P_r}{\sum_{a'} b_{ra'}} \left( \frac{b_{ra}}{N_{ra}} - \frac{b_{r-a}}{N_{r-a}} \right).$$

Consider flight  $i$  belonging to airline  $a$  and flight  $j$  belonging to airline  $-a$ . Flight  $i$  will be delayed by one slot due to flight  $j$  if  $\delta_{rij} = \alpha_{rai} - \alpha_{raj} > \tau_{ra}(\mathbf{b}_r)$ . As  $\alpha_{raj}$  and  $\alpha_{rai}$  are random variables,  $\delta_{rij}$  is as well. Let  $F$  be its CDF. The following lemma shows that we can restrict our attention to  $F$  and  $\tau$ .

LEMMA 5.  $u_{rai}(\mathbf{b}_r)$  is continuous, strictly concave, and strictly increasing in  $b_{ra}$  if and only if  $F(\tau_{ra}(\mathbf{b}_r))$  is.

$$\begin{aligned}
 u_{rai}(\mathbf{b}_r) &= \mathbb{E}_{\alpha_r} [V - \lambda_r d_{rai}] \\
 &= V - \lambda_r \mathbb{E}_{\alpha_r} \left[ -\alpha_{rai} + \sum_{j=1}^{N_{ra}} z_r \mathbb{I}(\alpha_{rai} > \alpha_{raj}) + \sum_{j=1}^{N_{r-a}} z_r \mathbb{I}(\delta_{rij} > p_{rai} - p_{r-aj}) \right] \\
 &= V + \lambda_r \frac{T_r}{2} - \lambda_r z_r 0.5(N_{ra} - 1) - \lambda_r z_r \sum_{j=1}^{N_{r-a}} \mathbb{E}_{\alpha_r} [\mathbb{I}(\delta_{rij} \geq p_{rai} - p_{r-aj})] \\
 &= V + \lambda_r \frac{T_r}{2} - \lambda_r z_r 0.5(N_{ra} - 1) - \lambda_r z_r \sum_{j=1}^{N_{r-a}} \mathbb{E}_{\alpha_r} [\mathbb{I}(\delta_{rij} \geq \tau_{ra}(\mathbf{b}_r))] \\
 &= V + \lambda_r \frac{T_r}{2} - \lambda_r z_r 0.5(N_{ra} - 1) - \lambda_r z_r \sum_{j=1}^{N_{r-a}} \mathbb{P}(\delta_{rij} \geq \tau_{ra}(\mathbf{b}_r)) \\
 &= V + \lambda_r \frac{T_r}{2} - \lambda_r z_r 0.5(N_{ra} - 1) - \lambda_r z_r N_{r-a} (1 - F(\tau_{ra}(\mathbf{b}_r))).
 \end{aligned}$$

All but the final term is independent of  $\mathbf{b}_r$ , so  $u_{rai}(\mathbf{b}_r)$  is strictly concave and strictly increasing iff the final term is.

We next consider the behavior of  $\tau$ .

LEMMA 6.  $\tau_{ra}(\mathbf{b}_r)$  is a continuous, strictly concave, and strictly increasing function of  $b_{ra}$ .

$$\begin{aligned}
 \frac{\partial}{\partial b_{ra}} \tau_{ra}(\mathbf{b}_r) &= \frac{\partial}{\partial b_{ra}} \left[ \frac{P_r}{\sum_{a'} b_{ra'}} \left( \frac{b_{ra}}{N_{ra}} - \frac{b_{r-a}}{N_{r-a}} \right) \right] \\
 &= P_r \frac{\left( \frac{\sum_{a'} b_{ra'}}{N_{ra}} \right) - \left( \frac{b_{ra}}{N_{ra}} - \frac{b_{r-a}}{N_{r-a}} \right)}{\left( \sum_{a'} b_{ra'} \right)^2} \\
 &= P_r \frac{b_{r-a} \left( \frac{1}{N_{ra}} + \frac{1}{N_{r-a}} \right)}{\left( \sum_{a'} b_{ra'} \right)^2} \tag{4}
 \end{aligned}$$

$$\begin{aligned}
 \frac{\partial^2}{\partial b_{ra}^2} \tau_{ra}(\mathbf{b}_r) &= \frac{\partial}{\partial b_{ra}} \left[ P_r \frac{b_{r-a} \left( \frac{1}{N_{ra}} + \frac{1}{N_{r-a}} \right)}{\left( \sum_{a'} b_{ra'} \right)^2} \right] \\
 &= P_r \frac{-2b_{r-a} \left( \frac{1}{N_{ra}} + \frac{1}{N_{r-a}} \right)}{\left( \sum_{a'} b_{ra'} \right)^3}
 \end{aligned}$$

Note that  $b_{r-a} \left( \frac{1}{N_{ra}} + \frac{1}{N_{r-a}} \right) > 0$  under the assumption that there is a minimum bid of  $\epsilon > 0$ . Thus, we have that  $\frac{\partial}{\partial b_{ra}} \tau_{ra}(\mathbf{b}_r) > 0$  and  $\frac{\partial^2}{\partial b_{ra}^2} \tau_{ra}(\mathbf{b}_r) < 0$ , as desired.

We now provide a closed form for  $F$  and its derivatives.

LEMMA 7.

$$F(\tau) = \begin{cases} \frac{1}{2T_r^2} (T_r^2 + 2T_r\tau - \tau^2) & \tau \geq 0 \\ \frac{(T_r + \tau)^2}{2T_r^2} & \tau < 0 \end{cases}$$

Thus its derivatives are

$$f(\tau) = \begin{cases} \frac{T_r - \tau}{T_r^2} & \tau \geq 0 \\ \frac{T_r + \tau}{T_r^2} & \tau < 0 \end{cases}$$

and

$$f'(\tau) = \begin{cases} -\frac{1}{T_r} & \tau > 0 \\ \frac{1}{T_r} & \tau < 0 \\ \text{undefined} & \tau = 0 \end{cases}$$

$$\begin{aligned} F(\tau) &= \mathbb{P}[\alpha_{ri} - \alpha_{rj} \leq \tau] \\ &= \int_{\max\{0, -\tau\}}^{T_r} \int_0^{\min\{T_r, y+\tau\}} \frac{1}{T^2} dx dy \end{aligned}$$

Thus, for  $\tau \geq 0$ , we have the CDF:

$$\begin{aligned} F(\tau) &= \int_0^{T_r - \tau} \int_0^{y+\tau} \frac{1}{T^2} dx dy + \int_{T_r - \tau}^{T_r} \int_0^{T_r} \frac{1}{T^2} dx dy \\ &= \frac{1}{T_r^2} \int_0^{T_r - \tau} (y + \tau) dy + \frac{1}{T_r} \int_{T_r - \tau}^{T_r} 1 dy \\ &= \frac{1}{T_r^2} (T_r^2 - \tau^2) + \frac{1}{T_r} \tau \\ &= \frac{1}{2T_r^2} (T_r^2 + 2T_r\tau - \tau^2) \end{aligned}$$

For  $\tau \leq 0$ , we have the CDF:

$$\begin{aligned} F(\tau) &= \int_{-\tau}^{T_r} \int_0^{y+\tau} \frac{1}{T^2} dx dy \\ &= \frac{1}{T^2} \int_{-\tau}^{T_r} (y + \tau) dy \\ &= \frac{(T_r + \tau)^2}{2T_r^2} \end{aligned}$$

The derivatives follow.

Finally, we characterize the behavior of  $F(\tau_{ra}(\mathbf{b}_r))$ .

LEMMA 8.  $F(\tau_{ra}(\mathbf{b}_r))$  is a continuous, strictly concave, and strictly increasing function of  $b_{ra}$  if  $P_r \leq \frac{T_r}{2} \min(N_{ra}, N_{r-a})$ .

To show that  $F(\tau_{ra}(\mathbf{b}_r))$  is strictly increasing with respect to  $b_{ra}$ , we have:

$$\frac{\partial}{\partial b_{ra}} F(\tau_{ra}(\mathbf{b}_r)) = f(\tau_{ra}(\mathbf{b}_r)) \frac{\partial}{\partial b_{ra}} \tau_{ra}(\mathbf{b}_r) \quad (5)$$

$f$  is the PDF of  $F$ , and so strictly positive. By Lemma 6,  $\frac{\partial}{\partial b_{ra}} \tau_{ra}(\mathbf{b}_r)$  is strictly positive. Thus, we have  $\frac{\partial}{\partial b_{ra}} F(\tau_{ra}(\mathbf{b}_r)) > 0$ , as desired.

Regarding the concavity of  $F(\tau_{ra}(\mathbf{b}_r))$  with respect to  $b_{ra}$ , we consider three cases: 1.  $\tau_{ra}(\mathbf{b}_r) > 0$ , 2.  $\tau_{ra}(\mathbf{b}_r) < 0$ , and 3.  $\tau_{ra}(\mathbf{b}_r) = 0$ . For  $\tau_{ra}(\mathbf{b}_r) \neq 0$ , we have:

$$\frac{\partial^2}{\partial b_{ra}^2} F(\tau_{ra}(\mathbf{b}_r)) = f(\tau_{ra}(\mathbf{b}_r)) \frac{\partial^2}{\partial b_{ra}^2} \tau_{ra}(\mathbf{b}_r) + f'(\tau_{ra}(\mathbf{b}_r)) \left( \frac{\partial}{\partial b_{ra}} \tau_{ra}(\mathbf{b}_r) \right)^2 \quad (6)$$

**Case 1:** For  $\tau > 0$ ,  $f(\tau) > 0$  and  $f'(\tau) < 0$ . Thus, when  $\tau_{ra}(\mathbf{b}_r) > 0$ , both terms are negative, so  $\frac{\partial^2}{\partial b_{ra}^2} F(\tau_{ra}(\mathbf{b}_r)) < 0$  as desired.

**Case 2:** For  $\tau < 0$ ,  $f(\tau) > 0$  and  $f'(\tau) > 0$ . Thus, when  $\tau_{ra}(\mathbf{b}_r) < 0$ , the term  $f'(\tau_{ra}(\mathbf{b}_r)) \left( \frac{\partial}{\partial b_{ra}} \tau_{ra}(\mathbf{b}_r) \right)^2$  is positive and a more careful analysis is needed. Applying Lemma 6 and Lemma 7, we have:

$$\begin{aligned} & f(\tau_{ra}(\mathbf{b}_r)) \frac{\partial^2}{\partial b_{ra}^2} \tau_{ra}(\mathbf{b}_r) + f'(\tau_{ra}(\mathbf{b}_r)) \left( \frac{\partial}{\partial b_{ra}} \tau_{ra}(\mathbf{b}_r) \right)^2 \\ &= \frac{1}{T_r^2} \left( T_r + \frac{P_r}{\sum_{a'} b_{ra'}} \left( \frac{b_{ra}}{N_{ra}} - \frac{b_{r-a}}{N_{r-a}} \right) \right) \left( -2P_r \frac{b_{r-a} \left( \frac{1}{N_{ra}} + \frac{1}{N_{r-a}} \right)}{\left( \sum_{a'} b_{ra'} \right)^3} \right) \\ &+ \frac{1}{T_r^2} \left( P_r \frac{b_{r-a} \left( \frac{1}{N_{ra}} + \frac{1}{N_{r-a}} \right)}{\left( \sum_{a'} b_{ra'} \right)^2} \right)^2 \\ &= \frac{P_r^2 b_{r-a} \left( \frac{1}{N_{ra}} + \frac{1}{N_{r-a}} \right)}{T_r^2 \left( \sum_{a'} b_{ra'} \right)^4} \left( \frac{-2T_r}{P_r} \sum_{a'} b_{ra'} - 2 \left( \frac{b_{ra}}{N_{ra}} - \frac{b_{r-a}}{N_{r-a}} \right) + b_{r-a} \left( \frac{1}{N_{ra}} + \frac{1}{N_{r-a}} \right) \right) \end{aligned}$$

Because the first two fractions are strictly positive,  $F(\tau_{ra}(\mathbf{b}_r))$  is strictly concave iff the remaining term is strictly less than zero.

$$\begin{aligned} & \frac{-2T_r}{P_r} \sum_{a'} b_{ra'} - 2 \left( \frac{b_{ra}}{N_{ra}} - \frac{b_{r-a}}{N_{r-a}} \right) + b_{r-a} \left( \frac{1}{N_{ra}} + \frac{1}{N_{r-a}} \right) \\ &= -2 \left( \frac{T_r}{P_r} + \frac{1}{N_{ra}} \right) b_{ra} + \left( \frac{-2T_r}{P_r} + \frac{1}{N_{ra}} + \frac{3}{N_{r-a}} \right) b_{r-a} < 0 \end{aligned}$$

Thus, when  $\tau_{ra}(\mathbf{b}_r) < 0$ , if  $\frac{2T_r}{P_r} \geq \frac{1}{N_{ra}} + \frac{3}{N_{r-a}}$ , then we have  $\frac{\partial^2}{\partial b_{ra}^2} F(\tau_{ra}(\mathbf{b}_r)) < 0$ , as desired. Thus,  $P_r \leq \frac{T_r}{2} \min(N_{ra}, N_{r-a})$  is a sufficient condition for strict concavity in this case.

**Case 3:**  $f$  is not differentiable at zero, so the preceding analysis cannot be used. However, as the function is continuous and strictly concave at all but perhaps a single point, this suffices to show it is strictly concave.

## Proof of Lemma 2

Consider an airport  $r$  and numbers of flights  $N_{r_1}$  and  $N_{r_2}$ . The total expected delay at this airport is  $\lambda_r z_r (N_{r_1} + N_{r_2})(N_{r_1} + N_{r_2} - 1)/2 - (N_{r_1} + N_{r_2})T_r/2$ . Since our simple model assumes the cost of delay is simply the amount incurred, and the value of each flight,  $V$ , is independent of the bids, the sum of the airline utilities is also constant, equal to  $(N_{r_1} + N_{r_2})V$  minus the expected delay. Thus, the second stage game is constant sum.

Now, consider the function  $v_1(\mathbf{b})$  that determines the utility of airline 1 given bids  $b$  at all airports (i.e., the sum of the utility  $u_{r_1i}$  across airports and flights). Because of the constant sum nature of the game, minimizing this function is equivalent to maximizing  $v_2(\mathbf{b})$ , the utility of airline 2. By Lemma 1, the function  $v_1(\mathbf{b})$  is continuous in the vectors  $\mathbf{b}_1$  and  $\mathbf{b}_2$ , strictly concave in  $\mathbf{b}_1$ , and strictly convex in  $\mathbf{b}_2$ . Therefore, because our strategy sets are compact and convex, by Sion's Minimax Theorem (Sion 1958) there is a unique equilibrium payoff for airline 1,  $v_1^*$ , which determines the unique equilibrium payoff for airline 2,  $v_2^*$ .

Now suppose that both players make the same bid at airport  $r$  (i.e.  $b_{r_1} = b_{r_2}$ ). By equation (4) from the proof of Lemma 6,  $\frac{\partial}{\partial b_{r_1}} [\tau_{r_1}(\mathbf{b}_r)] = \frac{\partial}{\partial b_{r_2}} [\tau_{r_2}(\mathbf{b}_r)]$ . Further, note that  $\tau_{r_1} = -\tau_{r_2}$  and  $f(\tau) = f(-\tau)$ . This allows us to show that the marginal benefits of increasing their bids are the same for both airlines:

$$\begin{aligned}
\frac{\partial}{\partial b_{r_2}} \left[ \sum_{i=1}^{N_{r_2}} u_{r_2i}(\mathbf{b}_r) \right] &= \lambda_r z_r N_{r_1} N_{r_2} \frac{\partial}{\partial b_{r_2}} [F(\tau_{r_2}(\mathbf{b}_r))] \\
&= \lambda_r z_r N_{r_1} N_{r_2} f(\tau_{r_2}(\mathbf{b})) \frac{\partial}{\partial b_{r_2}} [\tau_{r_2}(\mathbf{b}_r)] \\
&= \lambda_r z_r N_{r_1} N_{r_2} f(-\tau_{r_1}(\mathbf{b})) \frac{\partial}{\partial b_{r_1}} [\tau_{r_1}(\mathbf{b}_r)] \\
&= \lambda_r z_r N_{r_1} N_{r_2} f(\tau_{r_1}(\mathbf{b})) \frac{\partial}{\partial b_{r_1}} [\tau_{r_1}(\mathbf{b}_r)] \\
&= \lambda_r z_r N_{r_1} N_{r_2} \frac{\partial}{\partial b_{r_1}} [F(\tau_{r_1}(\mathbf{b}_r))] \\
&= \frac{\partial}{\partial b_{r_1}} \left[ \sum_{i=1}^{N_{r_1}} u_{r_1i}(\mathbf{b}_r) \right]
\end{aligned}$$

Now, we note that the best response for airline 1 can be written as a convex optimization problem.

$$\begin{aligned}
&\underset{\mathbf{b}_1}{\text{maximize}} && v_1(\mathbf{b}) \\
&\text{subject to} && \sum_r b_{r_1} = 1, \\
&&& b_{r_1} \geq \epsilon, \quad \forall r.
\end{aligned}$$

The Lagrangian dual of this convex optimization problem can be expressed as:

$$\begin{aligned}\mathcal{L}_D(\mu, \nu) &= \max_{\mathbf{b}_1} \left\{ v_1(\mathbf{b}) + \sum_r \mu_r (b_{r1} - \epsilon) + \nu (1 - \sum_r b_{r1}) \right\} \\ &= \sum_r \max_{b_{r1}} \{ N_{r1} u_{r1i}(\mathbf{b}_r) + \mu_r (b_{r1} - \epsilon) - \nu b_{r1} \} + \nu\end{aligned}$$

Thus, the first-order conditions for the inner maximization suggest that at each airport  $r$ , the following must hold for optimality:

$$\frac{\partial}{\partial b_{r1}} N_{r1} u_{r1i} + \mu_r = \nu$$

And, in particular, if  $\mu_r = 0$ , then  $\nu = \frac{\partial}{\partial b_{r1}} N_{r1} u_{r1i}$ . Combining this with the earlier insight that airlines face the same marginal utility when bids at an airport are equal allows us to construct an equilibrium solution. For each airport  $r$ , we consider two cases: 1.  $N_{r1} \leq N_{r2}$  and 2.  $N_{r1} > N_{r2}$ .

**Case 1:** Note that when  $N_{r1} \leq N_{r2}$ , if  $b_{r1} = b_{r2}$ , then  $\tau_{r1}(\mathbf{b}_r) \geq 0$ . Thus, we consider the marginal utility of airline 1 at airport  $r$  with respect to  $b_{r1}$  when  $\tau_{r1} \geq 0$ .

$$\frac{\partial}{\partial b_{r1}} N_{r1} u_{r1i} = \frac{\lambda_r z_r P_r b_{r1} (N_{r1} + N_{r2})}{(b_{r1} + b_{r2})^3 N_{r1} N_{r2} T_r^2} (b_{r1} N_{r2} (-P_r + N_{r1} T_r) + b_{r2} N_{r1} (P_r + N_{r2} T_r))$$

Evaluating this at  $b_{r2} = b_{r1}$ , we have:

$$\left[ \frac{\partial}{\partial b_{r1}} N_{r1} u_{r1i} \right]_{b_{r2}=b_{r1}} = \frac{\lambda_r z_r (N_{r1} + N_{r2}) P_r}{8 b_{r1} N_{r1} N_{r2} T_r^2} (P_r (N_{r1} - N_{r2}) + 2 N_{r1} N_{r2} T_r)$$

Given a dual multiplier  $\nu$ , we can find a dual feasible bidding strategy for both airlines by setting:

$$\begin{aligned}\left[ \frac{\partial}{\partial b_{r1}} N_{r1} u_{r1i} \right]_{b_{r2}=b_{r1}} &= \nu \implies \\ \bar{b}_{r1}(\nu) = \bar{b}_{r2}(\nu) &= \frac{\lambda_r z_r (N_{r1} + N_{r2}) P_r}{8 \nu N_{r1} N_{r2} T_r^2} (P_r (N_{r1} - N_{r2}) + 2 N_{r1} N_{r2} T_r)\end{aligned}$$

**Case 2:** When  $N_{r1} > N_{r2}$ , following the same reasoning as above leads us to a dual feasible bidding strategy for both airlines of:

$$\bar{b}_{r1}(\nu) = \bar{b}_{r2}(\nu) = \frac{\lambda_r z_r (N_{r1} + N_{r2}) P_r}{8 \nu N_{r1} N_{r2} T_r^2} (P_r (N_{r2} - N_{r1}) + 2 N_{r1} N_{r2} T_r)$$

We note that in each case, the corresponding bids are strictly positive assuming  $P_r \leq 2T_r \min(N_{ra}, N_{r-a})$ . This condition is satisfied by the statement of the lemma, so as long as  $\epsilon$  is less than the minimum constructed bid, each of these bids will exceed  $\epsilon > 0$ . Thus, we set  $\mu^* = 0$ ,  $\nu^* = \sum_r \bar{b}_{r1}(1)$ , and  $\mathbf{b}^* = \bar{\mathbf{b}}(\nu^*)$ . Under this choice of values, we note that the Karush Kuhn-Tucker conditions are satisfied, that is 1. the solution is stationary, because  $\nu^* = \frac{\partial}{\partial b_{r1}} N_{r1} u_{r1i}(\mathbf{b}_r^*)$  for all  $r$ ; 2.  $\mathbf{b}_1^*$  is primal feasible; 3.  $(\mu^*, \nu^*)$  is dual feasible; and 4. complementary slackness holds (i.e.,  $\mu_r^* b_{r1}^* = 0 \forall r$ ). This implies that the constructed bidding strategy is a Nash equilibrium, and therefore that it achieves the unique equilibrium payoff for the game.

### Proof of Lemma 3

By Lemma 2, the equilibrium allocation of priority at an airport is independent of the number of flights scheduled at other airports. Therefore, it suffices to show that, at a single airport, the optimal number of flights for an airline to schedule is nonincreasing in the number of flights the other airline schedules. The value of a flight is independent of the number of flights the other airline schedules. Thus, it suffices to show that the marginal delay from the adding a single flight is increasing in the number of the other airlines flights. The total delay for airline  $a$  with  $N_{ra}$  and  $N_{r-a}$  flights scheduled is

$$D_{ra}(N_{ra}, N_{r-a}) = N_{ra} \left( \lambda_r z_r \frac{N_{ra} - 1}{2} + \lambda_r z_r N_{r-a} \left( 1 - F \left( \frac{P_r}{2} \left( \frac{1}{N_{ra}} - \frac{1}{N_{r-a}} \right) \right) \right) - \frac{T_r}{2} \right)$$

While the  $N_{r-a}$  must be an integer, treating it as real-valued allows us to write

$$\begin{aligned} & \frac{\partial}{\partial N_{r-a}} [D_{ra}(N_{ra} + 1, N_{r-a}) - D_{ra}(N_{ra}, N_{r-a})] \\ &= \begin{cases} \lambda_r z_r \frac{N_{ra}^2 (4T_r^2 N_{r-a}^2 - P_r^2) + N_{ra} (4T_r^2 N_{r-a}^2 - P_r^2) - P_r^2 N_{r-a}^2}{8N_{ra} (1+N_{ra}) N_{r-a}^2 T_r^2} & N_{ra} \geq N_{r-a} \\ \lambda_r z_r \frac{N_{ra}^2 (4T_r^2 N_{r-a}^2 + P_r^2) + N_{ra} (4T_r^2 N_{r-a}^2 + P_r^2) + P_r^2 N_{r-a}^2}{8N_{ra} (1+N_{ra}) N_{r-a}^2 T_r^2} & N_{ra} < N_{r-a} \end{cases} \end{aligned}$$

The second case is strictly positive. As  $N_{ra} \geq 1$  and  $P_r^2 N_{r-a}^2 > 0$ , the first case is non-negative when

$$2(4T_r^2 N_{r-a}^2 - P_r^2) - P_r^2 N_{r-a}^2 \geq 0.$$

As we also have  $N_{r-a} \geq 1$ , it suffices that

$$P_r \leq \sqrt{8/3} T_r.$$

### Proof of Lemma 4

As before, it suffices to analyze a single airport in isolation. We begin by showing that, for a fixed total number of flights, the marginal cost that an airline faces is strictly increasing in the number of the flights it has scheduled.

LEMMA 9. *If  $P_r < \sqrt{6} T_r$ ,*

$$D_{ra}(N_{ra} + 2, N_{r-a}) - D_{ra}(N_{ra} + 1, N_{r-a}) > D_{ra}(N_{ra} + 1, N_{r-a} + 1) - D_{ra}(N_{ra}, N_{r-a} + 1).$$

We consider 4 cases, determined by the relationship between  $N_{ra}$  and  $N_{r-a}$  (which in turn determines the form of  $D$ ).

**Case 1:**  $N_{ra} \leq N_{r-a} - 2$ . After simplification we have

$$\begin{aligned} & D_{ra}(N_{ra} + 2, N_{r-a}) - D_{ra}(N_{ra} + 1, N_{r-a}) - D_{ra}(N_{ra} + 1, N_{r-a} + 1) + D_{ra}(N_{ra}, N_{r-a} + 1) \\ &= \frac{\lambda_r z_r}{8T_r^2 N_{ra} (1+N_{ra}) (2+N_{ra}) N_{r-a} (1+N_{r-a})} \times \\ & \quad \left( T_r^2 N_{r-a} (24N_{ra} + 36N_{ra}^2 + 12N_{ra}^3) (1+N_{r-a}) \right. \\ & \quad \left. + P_r^2 (2N_{ra} + 3N_{ra}^2 + N_{ra}^3 + 2N_{r-a} + 4N_{r-a}^2 + 2N_{r-a}^3 + N_{ra} N_{r-a} (1+N_{r-a})) \right), \end{aligned}$$

which is strictly positive.

**Case 2:**  $N_{ra} \geq N_{r-a} + 1$ . After simplification we have

$$\begin{aligned} & D_{ra}(N_{ra} + 2, N_{r-a}) - D_{ra}(N_{ra} + 1, N_{r-a}) - D_{ra}(N_{ra} + 1, N_{r-a} + 1) + D_{ra}(N_{ra}, N_{r-a} + 1) \\ &= \frac{\lambda_r z_r}{8T_r^2 N_{ra} (1 + N_{ra})(2 + N_{ra})N_{r-a}(1 + N_{r-a})} \times \\ & \quad \left( T_r^2 N_{r-a} (24N_{ra} + 36N_{ra}^2 + 12N_{ra}^3)(1 + N_{r-a}) \right. \\ & \quad \left. - P_r^2 (2N_{ra} + 3N_{ra}^2 + N_{ra}^3 + 2N_{r-a} + 4N_{r-a}^2 + 2N_{r-a}^3 + N_{ra}N_{r-a}(1 + N_{r-a})) \right). \end{aligned}$$

Using  $N_{r-a} \geq 1$  and gathering terms of similar orders, this is weakly positive as long as

$$48T_r^2 N_{ra} \geq P_r^2 (2N_{ra} + 2N_{r-a}),$$

$$72T_r^2 N_{ra}^2 \geq P_r^2 (3N_{ra}^2 + 4N_{r-a} + N_{ra}N_{r-a}), \text{ and}$$

$$24T_r^2 N_{ra}^3 \geq P_r^2 (N_{ra}^2 + 2N_{r-a} + N_{ra}N_{r-a}^2).$$

Thus,  $P_r < \sqrt{6}T_r$  suffices for to make it strictly positive in this case.

**Case 3:**  $N_{ra} = N_{r-a}$ . After simplification we have

$$\begin{aligned} & D_{ra}(N_{ra} + 2, N_{r-a}) - D_{ra}(N_{ra} + 1, N_{r-a}) - D_{ra}(N_{ra} + 1, N_{r-a} + 1) + D_{ra}(N_{ra}, N_{r-a} + 1) \\ &= \frac{6T_r^2 (2 + 3N_{ra} + N_{ra}^2) - P_r^2 (2 + 3N_{ra} + N_{ra}^2)}{T_r^2 (N_{ra} + 1)^2 (N_{ra} + 2)^2}. \end{aligned}$$

Again,  $P_r < \sqrt{6}T_r$  suffices to make it strictly positive.

**Case 4:**  $N_{ra} = N_{r-a} - 1$ . After simplification we have

$$\begin{aligned} & D_{ra}(N_{ra} + 2, N_{r-a}) - D_{ra}(N_{ra} + 1, N_{r-a}) - D_{ra}(N_{ra} + 1, N_{r-a} + 1) + D_{ra}(N_{ra}, N_{r-a} + 1) \\ &= \lambda_r z_r \left( 1.5 + \frac{P_r^2}{4T_r^2 N_{ra} (N_{ra} + 1)} \right), \end{aligned}$$

which is strictly positive.

We now show that all equilibria have the same number of flights, except for a set of values of  $V$  of measure zero. Suppose that there are equilibria  $(N_{r1}, N_{r2})$  and  $(N'_{r1}, N'_{r2})$  with  $N_{r1} + N_{r2} > N'_{r1} + N'_{r2}$ . Consider such a pair where some airline schedules the same number of flights in both equilibria (WLOG let  $N_{r2} = N'_{r2}$ ). Both strategies are equilibria and marginal costs are strictly increasing, so  $N_{r1} = N'_{r1} + 1$  and airline 1 is indifferent between them. For each pair  $(N_{r1}, N_{r2})$  and corresponding equilibrium bidding strategy, the marginal cost of scheduling another flight is fixed and independent of  $V$ , thus there is exactly one  $V$  (equal to the marginal cost) which results in indifference towards scheduling an additional flight. Thus, the set of  $V$  for which this can occur has the same cardinality as  $\mathbb{Z}$  and therefore has measure zero in  $\mathbb{R}$ .

Otherwise, both airlines schedule different numbers of flights in each equilibrium and we derive a contradiction. Since equilibria are symmetric, WLOG let airline 1 have the most flights in both equilibria (i.e.  $N_{r1} \geq N_{r2}$  and  $N'_{r1} \geq N'_{r2}$ ). Because marginal costs are strictly increasing, we cannot have both  $N_{r1} > N'_{r1}$  and  $N_{r2} > N'_{r2}$ , but one of these must hold. Suppose only the former holds. Then  $N'_{r1} + N'_{r2} - N_{r2} < N_{r1}$ . Let  $\Delta_{ra}(N_{r1}, N_{r2}) = D_{ra}(N_{r1} + 1, N_{r2}) - D_{ra}(N_{r1}, N_{r2})$ . As  $(N'_{r1}, N'_{r2})$  is an equilibrium, airline 1 must be at best indifferent about adding another flight (i.e.  $\Delta_{r1}(N'_{r1}, N'_{r2}) \geq V$ ). As  $\Delta_{ra}(\cdot, N_{r-a})$  is increasing and  $(N_{r1}, N_{r2})$  is an equilibrium, airline 1 must be at worst indifferent about adding another flight when it has less than  $N_{r1}$ . In particular,  $\Delta_{r1}(N'_{r1} + N'_{r2} - N_{r2}, N_{r2}) \leq V$ . By Lemma 9 we have

$$V \geq \Delta_{r1}(N'_{r1} + N'_{r2} - N_{r2}, N_{r2}) > \Delta_{r1}(N'_{r1}, N'_{r2}) \geq V,$$

which is a contradiction. Suppose instead only the latter holds. Then  $N'_{r2} + N'_{r1} - N_{r1} < N_{r2}$ . Similar to the previous case,  $\Delta_{r2}(N'_{r1}, N'_{r2}) \geq V$  and  $\Delta_{r2}(N_{r1}, N'_{r2} + N'_{r1} - N_{r1}) \leq V$ . Again by Lemma 9 we have

$$V \geq \Delta_{r2}(N_{r1}, N'_{r2} + N'_{r1} - N_{r1}) > \Delta_{r2}(N'_{r1}, N'_{r2}) \geq V,$$

a contradiction. Thus, all equilibria have the same number of flights.

It remains to show that the equilibrium number of flights is decreasing in the size of the pool. To do so we make use of the following lemma that shows that, for both airlines, marginal costs are increasing in  $P_r$ .

LEMMA 10. *If  $P_r \leq 2T_r N_{r-a}$ ,  $\frac{d}{dP_r} [D_{ra}(N_{ra} + 1, N_{r-a}) - D_{ra}(N_{ra}, N_{r-a})] > 0$ .*

$$\begin{aligned} & \frac{d}{dP_r} [D_{ra}(N_{ra} + 1, N_{r-a}) - D_{ra}(N_{ra}, N_{r-a})] \\ &= \begin{cases} \lambda_r z_r \frac{N_{ra}^2(2T_r N_{r-a} + P_r) + N_{ra}(2T_r N_{r-a} + P_r) - P_r N_{r-a}^2}{4N_{ra}(1+N_{ra})N_{r-a}T_r^2} & N_{ra} \geq N_{r-a} \\ \lambda_r z_r \frac{N_{ra}^2(2T_r N_{r-a} - P_r) + N_{ra}(2T_r N_{r-a} - P_r) + P_r N_{r-a}^2}{4N_{ra}(1+N_{ra})N_{r-a}T_r^2} & N_{ra} < N_{r-a} \end{cases} \end{aligned}$$

The first case is always strictly positive. The second case is strictly positive as long as  $P_r \leq 2T_r N_{r-a}$ .

Since marginal costs only increase with  $P_r$ , we can use Lemma 9 in the same manner as before to contradict the possibility that the equilibrium number of flights has increased.

## References

Ball, M. O., G. Donohue, K. Hoffman. 2006. Auctions for the safe, efficient, and equitable allocation of airspace system resources. P. Cranton, Y. Shoham, R. Steinberg, eds., *Combinatorial Auctions*. MIT Press, Cambridge, MA, 507–538.

- Ball, M. O., R. Hoffman, A. Mukherjee. 2010. Ground delay program planning under uncertainty based on the ration-by-distance principle. *Transportation Sci.* **44**(1) 1–14.
- Brueckner, J. K. 2002. Airport congestion when carriers have market power. *Am. Econ. Rev.* **92**(5) 1357–1375.
- Brueckner, J. K. 2005. Internalization of airport congestion: a network analysis. *Int. J. Ind. Organ.* **23** 599–614.
- Chang, K., K. Howard, R. Oiesen, L. Shishler, M. Tanino, M. Wambsganns. 2001. Enhancements to the faa ground-delay program under collaborative decision-making. *Interfaces* **31**(1) 57–76.
- Fearing, D., C. Barnhart. 2011. Evaluating air traffic flow management in a collaborative decision-making environment. *Transportation Res. Rec. Special issue on Public Sector Aviation, Graduate Research Awards, 2009 - 2010*(2206) 10 – 18.
- Harsha, P. 2008. Mitigating airport congestion: market mechanisms and airline response models. Ph.D. thesis, Massachusetts Institute of Technology.
- Hoffman, K., F. Berardino, R. Golaszewski, L. Sherry, K. Aqbal, G. Hunter, D. Escala, H. Gao. 2011. Results of a congestion management study for day-of-flight. Presentation at INFORMS TSL Workshop on Congestion Management of Transportation Systems on the Ground and in the Air, Asilomar, CA. Available at [http://www.informs.org/content/download/244265/2317985/file/Hoffman-Asimalar\\_presentation\\_V2.pdf](http://www.informs.org/content/download/244265/2317985/file/Hoffman-Asimalar_presentation_V2.pdf).
- Kelly, F. P., A. K. Maulloo, D. Tan. 1998. Rate control in communication networks: Shadow prices, proportional fairness, and stability. *J. of the Operational Research Society* **49** 237–252.
- Manley, B. 2008. Minimizing the pain in air transportation: analysis of performance and equity in ground delay programs. Ph.D. thesis, George Mason University.
- Milgrom, P., J. Roberts. 1990. Rationalizability, learning, and equilibrium in games with strategic complement- arities. *Econometrica* **58**(6) 1255–1277.
- Morrison, S. A., C. Winston. 2007. Another look at airport congestion pricing. *Am. Econ. Rev.* **97**(5) 1970–1977.
- Rosen, J. B. 1965. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica* **33**(3) 520–534.
- Sion, Maurice. 1958. Resource demand and supply in BitTorrent content-sharing communities. *Pacific J. Math.* **8**(1) 171–176.
- Tarski, A. 1955. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics* **5** 285–309.
- Topkis, D. M. 1979. Equilibrium points in nonzero-sum n-person submodular games. *Siam Journal of Control and Optimization* **17** 773–787.

United States Department of Transportation. May 13, 2009. Transportation secretary LaHood proposes to cancel slot auctions for New York-area airports, announces discussions on reducing congestion. Press Release DOT 66-09. Available at <http://www.dot.gov/affairs/2009/dot6609.htm>.

Vazirani, V. 2011. A market for air traffic flow management. Working paper.

Vossen, T., M. O. Ball. 2006. Slot trading opportunities in collaborative ground delay programs. *Transportation Sci.* **40**(1) 29–43.