University of Pennsylvania
ScholarlyCommons

Operations, Information and Decisions Papers

Wharton Faculty Research

5-2005

Revenue Management Games: Horizontal and Vertical Competition

Serguei Netessine
University of Pennsylvania

Robert. A. Shumsky

Follow this and additional works at: https://repository.upenn.edu/oid_papers

Part of the [Business Administration, Management, and Operations Commons](#), and the [Strategic Management Policy Commons](#)

Recommended Citation

Netessine, S., & Shumsky, R. A. (2005). Revenue Management Games: Horizontal and Vertical Competition. *Management Science*, 51 (5), 813-831. <http://dx.doi.org/10.1287/mnsc.1040.0356>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/oid_papers/199
For more information, please contact repository@pobox.upenn.edu.

Revenue Management Games: Horizontal and Vertical Competition

Abstract

A well-studied problem in the literature on airline revenue (or yield) management is the optimal allocation of seat inventory among fare classes, given a demand distribution for each class. In practice, the seat allocation decisions of one airline affect the passenger demands for seats on other airlines. In this paper, we examine the seat inventory control problem under both horizontal competition (two airlines compete for passengers on the same flight leg) and vertical competition (different airlines fly different legs on a multileg itinerary). Such vertical competition can be the outcome of a code-sharing agreement between airlines, because each airline sells seats on the partner airlines' flights but the airlines are unwilling, or unable, to coordinate yield management decisions. We provide a general sufficient condition under which a pure-strategy Nash equilibrium exists in these revenue management games, and we also compare the total number of seats available in each fare class with, and without, competition. Analytical results as well as numerical examples demonstrate that more seats are protected for higher-fare passengers under horizontal competition than when a single airline acts as a monopoly. Under vertical competition the booking limit may be higher or lower, however, than the monopoly level, depending on the demand for connecting flights in each fare class. Finally, we discuss revenue-sharing contracts that coordinate the actions of both airlines.

Keywords

revenue management, yield management, competition, nash equilibrium, airlines

Disciplines

Business Administration, Management, and Operations | Strategic Management Policy

Revenue Management Games: Horizontal and Vertical Competition

Serguei Netessine
The Wharton School
University of Pennsylvania
Philadelphia, PA 19103-6340
netessine@wharton.upenn.edu

Robert A. Shumsky
Simon School of Business Administration
University of Rochester
Rochester, NY 14627
shumsky@simon.rochester.edu

Published as:

Netessine, S. and R. A. Shumsky. 2005. "Revenue Management Games," *Management Science* 51(5), pp. 813–831.

Revenue Management Games: Horizontal and Vertical Competition

Abstract: A well-studied problem in the literature on airline revenue (or yield) management is the optimal allocation of seat inventory among fare classes given a demand distribution for each class. In practice, the seat allocation decisions of one airline affect the passenger demands for seats on other airlines. In this paper we examine the seat inventory control problem under both horizontal competition (two airlines compete for passengers on the same flight leg) and vertical competition (different airlines fly different legs on a multi-leg itinerary). Such vertical competition can be the outcome of a code sharing agreement between airlines, for each airline sells seats on the partner airlines' flights but the airlines are unwilling, or unable, to coordinate yield management decisions. We provide a quite general sufficient condition under which a pure-strategy Nash equilibrium exists in these 'revenue management games', and we also compare the total number of seats available in each fare class with, and without, competition. Analytical results as well as numerical examples demonstrate that under horizontal competition more seats are protected for higher-fare passengers than when a single airline acts as a monopoly, while under vertical competition the booking limit may be higher, or lower, than the monopoly level, depending upon the demand for connecting flights in each fare class. Finally, we discuss revenue-sharing contracts that coordinate the actions of both airlines.

1. Introduction

Consider an airline customer looking for an early-morning flight from Rochester, NY to Chicago in May 2003. The traveler can choose between two airlines, American and United, who offer flights at nearly identical times (6:00 am and 6:10 am, respectively) at identical prices (both charge \$266 for 14-day and \$315 for 7-day advance purchase round-trip tickets). Now suppose that our customer wishes to purchase American's 14-day advance ticket. If the seats allocated to the 14-day fare class have sold out, it is likely that the customer will attempt to purchase a ticket in the same fare class on the United flight that departs 10 minutes later. In general, the allocation of seat inventory among fare classes by one airline affects the quantity of customer demand, and optimal seat allocation, of the other airline. If the airlines are competitors and do not collaborate on seat allocation decisions, then the decisions that arise out of the resulting game can differ significantly from the seat allocations that would be optimal for a single decision-maker with control over both airlines.

Here we use the term 'horizontal competition' when describing airlines that compete for customers on a single flight-leg. We will use the term 'vertical competition' to refer to different airlines that fly different legs on a multi-leg itinerary. For example, a traveler who wishes to fly from Cleveland to Tokyo can purchase a convenient two-leg itinerary with the first leg operated by Continental Airlines from Cleveland to Los Angeles, and the second leg operated by Northwest Airlines from Los Angeles to Tokyo. (We call this customer a 'connecting passenger.')

Because these two airlines have formed an alliance that allows each airline to sell tickets for flights operated by the alliance partner (an arrangement called *code sharing*), the Cleveland-Los Angeles-Tokyo ticket can be purchased from either airline. In this case, the term 'vertical competition' may seem to be a misnomer, for airlines in such an alliance cooperate by code sharing. However, technical and legal barriers often prevent the airlines from explicitly coordinating both pricing and seat inventory decisions.¹ As was true for horizontal competition, the seat inventory decisions of one airline affect customer demand and inventory decisions of the other, and the absence of perfect coordination can lead to inventory allocation decisions that differ significantly from the system-optimal solution.

In this paper we examine how both horizontal and vertical competition affect airline seat inventory decisions, and we consider how airlines in an alliance may coordinate these decisions by agreements

¹ Boyd (1998a) describes how coordinating the revenue management systems of multiple airlines can be a "logistical nightmare", and he explains that an airline may believe that its own proprietary system provides a competitive advantage and therefore the airline may resist merging systems with an ally. In addition, many alliances are not immunized from anti-trust regulations, and airline revenue managers for these airlines are concerned that the Federal Government may see overt coordination of seat inventory decisions as a violation of antitrust law (Zuckerman, 2002).

similar to revenue-sharing contracts. There is evidence that over the past ten years alliances have become more numerous and increasingly important for the airline industry. The U.S. airlines have recently formed the two largest alliances in history (United Airlines and US Airways in 2002, and Delta Airlines joined Northwest and Continental in 2003). The first major international alliances were formed in the early and mid-1990's, and international passenger traffic between hubs of code-sharing partners increased by 43% between 1992 and 1997, while international traffic between other hubs increased by just 7% (U. S. Department of Transportation, 2000). There are many advantages to such alliances. They allow airlines to expand into new markets and to increase the frequency of flights offered within markets served by both airlines. If given antitrust immunity, alliance partners can increase profits by coordinating activities such as marketing, purchasing, luggage handling and flight scheduling. These partners may also apply revenue management techniques to optimize joint revenues (for a thorough summary of the benefits and costs of airline alliances, see Fernandez de la [Torre, 1999](#)). However, as we mentioned above, there can be significant barriers to such coordination, and the airlines struggle to find decentralized mechanisms to integrate revenue management decisions.

There is a substantial economics literature analyzing airline behavior under competition and the impact of airline alliances. There is also a recent stream of operations research literature on the problem of optimal seat allocation. However, to our knowledge there are no published papers that place the seat allocation problem in a competitive framework, examine the impact of airline alliances on these decisions, or specify contracts that enable airlines to coordinate decisions. We address these issues here, for an understanding of how competition affects revenue management decisions would be helpful for airline managers who negotiate contracts with potential alliance partners or design systems to implement alliances, as well as for government regulators who must evaluate the impact of airline alliances.

After a review of the literature in Section 2, in Section 3 we describe the traditional single-airline, single-leg, two-class yield management problem, and we introduce a condition on the demand distribution that will be useful in later sections. We discuss horizontal competition in Section 4. In this case, each airline is faced with an initial demand from passengers who wish to purchase tickets, but each airline may also sell tickets to passengers that were denied a reservation on the competing airline. In Section 5 we consider vertical competition: each airline has both local and connecting demand, but the number of connecting passengers booked on one airline is limited by the number booked by an alliance partner. In both cases the optimal capacity limits for each class (the booking limits) on each airline are interdependent. We compare the revenue management policies of competing airlines with the policy of a monopolist who operates all flights or, equivalently, the policies of airlines that cooperate to maximize total profits. Analytical results for special cases as well as numerical examples demonstrate that under

horizontal competition more seats are protected for higher-fare passengers than when a single airline acts as a monopoly, while under vertical competition the booking limit may be higher, or lower, than the monopoly level, depending upon the quantity of connecting demand in each fare class. In Section 6 we present revenue-sharing contracts that coordinate the actions of both airlines, and in Section 7 we discuss managerial insights and consider directions for further work.

2. Related Literature

Publications that consider the interactions among economic forces, strategic airline market entry decisions, and airline schedules include the network design models of Lederer and Nambimamdom (1999), Dobson and Lederer (1994), and the empirical work by [Borenstein and Rose \(1994\)](#). Another body of research focuses on the airlines' scheduling decisions under competition using variants of the spatial model developed by Hotelling (1929). See, for example, the recent empirical papers by Borenstein and [Netz \(1999\)](#) and [Richard \(2003\)](#). These papers focus on broad competitive problems and ignore the specifics of seat inventory allocation. In this paper we will not be concerned with the reasons airlines schedule their flights at the same time or with the pricing decision for each flight. Rather, we will concentrate on the implications of competitive scheduling on seat inventory control.

There are numerous papers in the area of revenue management that focus on airline yield management. For fundamental results on the general subject of seat inventory control see [Belobaba \(1989\)](#), [Brumelle et al. \(1990\)](#), and a useful literature review by McGill and van Ryzin (1999). Some of our results for horizontal competition are related to the unpublished work by Li and Oum (1998), which describes a seat allocation problem for two airlines in competition but uses a relatively restrictive assumption about how demand is allocated among airlines and identifies one, symmetric equilibrium. Zhao and Atkins (2002) describe a model with two airlines competing for passengers in one demand class. [Belobaba and Wilson \(1997\)](#) describe a simulation model that is used to evaluate the benefits of yield management systems for airlines under competition.

The literature on inventory management has seen a stream of closely related papers devoted to competition among firms in which the firms determine inventory levels and customers may switch among firms until a suitable product is found. This has been described as a 'newsvendor game' or as 'inventory competition' and is related to our concept of *horizontal competition*. [Parlar \(1988\)](#) examines the competition between two retailers facing independent demands. [Lippman and McCardle \(1997\)](#) examine both the two-firm game and a game with an arbitrary number of players. In their models, initial industry demand is allocated among the players according to a pre-specified 'splitting rule.' They establish the existence of a pure strategy Nash equilibrium and show that the equilibrium is unique when the initial

allocation is deterministic and strictly increasing in the total industry demand for each player. Recent extensions of these models include Mahajan and van Ryzin (1999) who model demand as a stochastic sequence of utility-maximizing customers. For an arbitrary number of firms, they demonstrate that an equilibrium exists and show that it is unique for a symmetric game. [Netessine and Rudi \(2003\)](#) analyze a problem similar to [Parlar \(1988\)](#) but for an arbitrary number of products. Given mild parametric assumptions they establish the existence of, and characterize, a unique, globally stable Nash equilibrium.

Many of these papers compare total inventory levels under competition with inventory levels when firms cooperate. [Lippman and McCardle \(1997\)](#) show that competition can lead to higher inventories, and Mahajan and van Ryzin (1999) derive similar results given their dynamic model of customer purchasing. On the other hand, with the substitution structure of the model of Netessine and [Rudi \(2003\)](#), under competition some firms may stock less than under centralization. It is not immediately clear how these results can be extended to our model of horizontal competition, for the problem described here involves the allocation of a fixed inventory between two customer classes. Therefore, ‘overstocking’ of inventory in one class would imply ‘understocking’ in the other class. We will also see that in our problem, an airline’s demand depends upon its own booking-limit decision, a complication not seen in any newsvendor game.

In addition, our demand model for each fare class is more general than those that have been used in the literature. Lippman and McCardle’s stylized splitting rules generate demands that are either independent or perfectly correlated, while Parlar only considers independent demands. In our model, customer demand may follow an arbitrary distribution. Besides being a more natural description of real-world demand, allowing a general demand distribution leads to additional insights. For example, the traditional newsvendor game always has a pure strategy equilibrium, but ours may not. We will see that the correlation structure of demand between products has an important impact on the dynamics of the game, and our model allows us to find general conditions under which the game has an equilibrium.

Analysis of *vertical competition* dates back to [Spengler \(1950\)](#) who demonstrated in a simple supply chain that the effect of double-marginalization leads to suboptimal system performance. Since then there have been numerous studies of vertical competition in the supply chain setting but, to the best of our knowledge, no analysis in the revenue management setting. While there is both theoretical and empirical research investigating the impact of code sharing agreements on ticket prices (see, for example, [Brueckner and Whalen 2000](#), and [Brueckner 2003](#)), to the best of our knowledge there is no published work investigating the inefficiencies that arise during day-to-day yield management decisions.

The specific terms of code sharing agreements are usually not available to public. However,

Wynne (1995) and Boyd (1998a) describe some of the coordination mechanisms that are used in practice. Besides short accounts in Talluri and van Ryzin (2004) and Boyd (1998b), there has been no published analytical analysis of contracts among airline alliance partners. In general, there is much confusion, and very little theoretical support, for alliance partners when they attempt to design contracts. The contracts we propose are related to work of Cachon and Lariviere (2000) on revenue sharing in supply chain management. See also [Cachon \(2004\)](#) for the survey of contracting literature in operations management.

3. The Single-Airline, Single-Leg Problem

Before discussing the effects of competition on yield management practices, we review the traditional stand-alone yield management problem, introduce some notation, and derive a result that will be used throughout the paper. Let p_k be the net revenue from a passenger in class $k=L,H$ (low-fare and high-fare), which takes into account the variable cost of flying that passenger. Demand for tickets at these prices is represented by the random variables D_k , $k = L, H$. A ticket purchased for either fare provides access to the same product: a coach-class seat on one flight leg. We assume that demand for low-fare tickets occurs before demand for high-fare tickets, which is the case when advance-purchase requirements are used to distinguish between customers with different valuations on price and flexibility. Customers who prefer a low fare and are willing to accept the purchase restrictions will be called 'low-fare customers'. Customers who prefer to purchase later, at the higher price, are called 'high-fare customers'. We assume that there are no customer cancellations.

To maximize expected profits, the airline establishes a booking limit B for low-fare seats. Note that the establishment of a booking limit is an optimal policy for each airline – see [Brumelle et al. \(1990\)](#). Once the booking limit is reached, the low fare is closed. Sales of high-fare tickets are accepted until either the airplane is full or the flight departs.

This model contains only two fare classes, when in reality there may be many more (see Belobaba 1998 for an introduction to the complexities of real-world yield management systems). We also assume that the airlines' booking limits are static. That is, the booking limit is set before demand is realized and no adjustments are made as low-fare demand is observed. This model simplifies other aspects of the actual environment. For example, we assume that a passenger denied a low-fare booking does not attempt to upgrade to the high-fare class and that a passenger, when first denied a ticket, will not shift to a later or earlier flight operated by the same airline. These assumptions appear often in the literature on non-competitive revenue management, and, given the added complexity of our problem involving competition, we adopt them here as well.

Given that the aircraft has capacity C , the airline's optimization problem is

$$\max_B \pi = E \left[p_L \min(D_L, B) + p_H \min(D_H, C - \min(D_L, B)) \right]. \quad (3.1)$$

The first-order condition is,

$$\begin{aligned} \frac{\partial \pi}{\partial B} &= p_L \Pr(D_L > B) - p_H \min(D_H > C - B, D_L > B) \\ &= \Pr(D_L > B) (p_L - p_H \min(D_H > C - B | D_L > B)) = 0. \end{aligned} \quad (3.2)$$

If D_L, D_H are independent, it is easy to see that there is a unique solution. To address dependent demands, Brumelle et al. (1990) introduced the “monotonic association property.” Given that $\Pr(D_H > C - B | D_L > B)$ is monotonically increasing in B (‘monotonic association’), there is a unique solution to (3.2). Here we introduce a related property, that D_L and D_H are ‘totally positive of order 2’ (TP₂). TP₂ implies that realizations of the random variables are more likely to be high together, or low together, than to be mixed ‘high’ and ‘low’ (see Joe 1997 for a thorough discussion of TP₂ and its properties).

Proposition 1. *Suppose D_L, D_H are TP₂ (Totally Positive of Order 2). Then there is a unique solution to the optimality condition (3.2) and the objective function (3.1) is quasi-concave.*

Proof: From Proposition 2.3 of Joe (1997), TP₂ implies that D_L and D_H are right-tail increasing, i.e., $\Pr(D_H > C - B | D_L > B)$ is increasing in B . The result follows. ■

In the proof, the sufficient condition on D_L, D_H is actually weaker than the TP₂ condition: the two random variables must be ‘right-tail increasing’ (RTI), a property similar to (but stronger than) monotonic association. However, working with the TP₂ property is convenient, for we know that there are many useful bivariate distributions that are TP₂. These include any set of independent random variables, the multivariate logistic, Gamma and F distributions, and the bivariate normal distribution with positive correlation ([Karlin and Rinott, 1980](#), and [Tong, 1990](#)). The TP₂ property can also be extended to the *multivariate* TP₂ property (MTP₂). Both TP₂ and MTP₂ will be useful in the next two sections.

4. Horizontal Competition

Suppose two airlines offer direct flights between the same origin and destination, with departures and arrivals at similar times. We use subscripts $i, j=1, 2$ to distinguish two competing airlines. Flights have the seat capacity C_i and there are two fare classes available for passengers: ‘low-fare’ and ‘high-fare.’ If either type of customer is denied a ticket at one airline, the customer will attempt to purchase a ticket from the

other airline and we call these “overflow passengers” (see Figure 1). Therefore, both airlines are faced with a random initial demand for each fare class as well as demand from customers who are denied tickets by the other airline. Passengers denied reservations by both airlines are lost. Figure 1 shows the overflow processes. We assume that both airlines’ prices satisfy $p_{Li} < p_{Hi}$. We also assume that the random variables D_{ki} have nonnegative support and that the cumulative distributions are differentiable.

Figure 1: Horizontal Competition

Note that all results presented in this section also apply to a model in which some fraction (less than one) of passengers denied a ticket on one airline attempt to purchase a ticket from the other airline, while some fraction (greater than zero) are lost to both airlines. To simplify the model and minimize the number of parameters, we assume that all passengers denied a ticket from their first choice overflow to their second-choice airline.

The model describes two airlines engaging in a non-cooperative game with complete information. Each airline attempts to maximize its profits by adjusting its booking levels. In other words, the booking level $B_i \in [0, C_i]$ is the strategy space of airline i (for simplicity, we assume that the booking level may be any real number in this range). Each airline knows the strategy spaces and demand distributions of its own flight as well as those of the competing airline. Throughout the paper we will use B_i^c to denote equilibrium booking limits set by competing airlines, and we will be comparing these with B_i^a , the optimal booking limits established by an alliance of two airlines that coordinate yield management decisions (or, a profit-maximizing monopoly).

An important assumption of the model is that the initial demands D_{ki} are exogenous; they are not affected by the booking limits chosen by each airline. This assumption is consistent with the newsvendor games of Parlar (1988) and Lippman and McCardle (1997). However, one might argue that the booking limits determine seat availability, and that in the long run this aspect of service quality affects initial demand. A more complete model would incorporate this relationship between booking limits and

demand, and the solution would supply equilibrium demands as well as equilibrium booking limits. For our application, however, the relationship between booking limits and demand is weakened by marketing efforts such as advertising and frequent-flyer programs. In addition, the use of travel agents and on-line reservation tools reduces the marginal search cost associated with making each booking. Given low search costs, the decision as to which airline to query first may depend on factors that dominate the likelihood that the query will result in a booking.

The order of events in the game is consistent with the traditional yield management problem:

1. Airlines establish booking limits B_1 and B_2 .
2. Low-fare passengers arrive to their first-choice airlines and are accommodated up to the booking limits. Low-fare passengers not accommodated on their first-choice airlines 'spill' to the alternate airlines and are accommodated up to the booking limits.
3. High-fare passengers arrive to their first-choice airlines and are accommodated with any remaining seats, up to capacity C_i in each aircraft. High-fare passengers not accommodated on their first-choice airlines 'spill' to the alternate airlines and are accommodated in any remaining seats, up to capacity C_i in each aircraft.

Note that for an airline to enforce a booking-limit policy it is not necessary to distinguish between 'original' and 'overflow' passengers. Once demand is realized, each airline simply observes the arrival of low-fare passengers and then the arrival of high-fare passengers. Each group of customers contains a mixture of first-choice and overflow passengers.

To describe the problem in terms of customer demand and booking limits, define:

$$D_{Li}^T = D_{Li} + (D_{Lj} - B_j)^+, \text{ total demand for low-fare tickets on airline } i, i=1, j=2 \text{ and } i=2, j=1.$$

$$R_i = C_i - \min(D_{Li}^T, B_i), \text{ the number of seats available for high-fare passengers on airline } i = 1, 2.$$

$$D_{Hi}^T = D_{Hi} + (D_{Hj} - R_j)^+, \text{ total demand for high-fare tickets, } i=1, j=2 \text{ and } i=2, j=1.$$

The total revenue for airline i is

$$\pi_i = E \left[p_{Li} \min(D_{Li}^T, B_i) + p_{Hi} \min(D_{Hi}^T, R_i) \right]. \quad (4.1)$$

Each airline will maximize this expression, given the booking limit of its competitor. It will be instructive to examine the first derivative of this objective function. It is tedious to find the derivative by the traditional methods (e.g., applying Leibnitz's rule). Instead, by applying the techniques described in Rudi (2000), we find for $i=1, j=2$ and $i=2, j=1$,

$$\begin{aligned} \frac{\partial \pi_i}{\partial B_i} = & p_{Li} \Pr(D_{Li}^T > B_i) - p_{Hi} \Pr(D_{Hi}^T > C_i - B_i, D_{Li}^T > B_i) \\ & - p_{Hi} \Pr(D_{Li} > B_i, D_{Lj}^T < B_j, D_{Hj} > R_j, D_{Hi}^T < C_i - B_i). \end{aligned} \quad (4.2)$$

Although this is a complex expression, there is a straightforward interpretation for each term. An incremental increase in the booking limit B_i by airline i has three effects on that airline's total revenue. First, revenue from low-fare customers increases with probability $\Pr(D_{Li}^T > B_i)$. Second, revenue from the high-fare customers decreases with probability $\Pr(D_{Hi}^T > C_i - B_i, D_{Li}^T > B_i)$. While these two effects are direct consequences of the change in B_i , there is a third, indirect effect. Revenue from high-fare customers may decrease due to the following sequence of events: (i) an increase in B_i may reduce the overflow of low-fare customers from i to j , (ii) a reduction in the number of low-fare customers at j may increase the number of seats available for high-fare customers at j , (iii) this may reduce the overflow of high-fare customers from j to i and (iv) a decline in the overflow from j may reduce the number of high-fare customers accommodated at i . The probability of this sequence of events is the third term on the right-hand side of equation (4.2), which implies that an increase in the booking limit of airline i can result in a decrease in high-fare demand to airline i . Hence, even though there is no explicit dependence between airline's booking limit and its demand, such dependence is introduced through the game. Because the strategy spaces of the airlines are compact and the payoff functions are continuous, a Nash equilibrium in mixed strategies must exist. However, a pure-strategy Nash equilibrium may or may not exist for airlines playing this game.

4.1 Nash Equilibrium Conditions

In the following proposition we show that if D_{Li} and D_{Hi}^T are TP₂ then the game has a pure strategy Nash equilibrium. In a corollary we show that these conditions also hold if $\underline{D} = (D_{L1}, D_{L2}, D_{H1}, D_{H2})$ is MTP₂. This result indicates that the existence of a pure strategy equilibrium depends upon the correlation of demands among fare classes *within* each airline, and does not directly depend upon the demand correlation *between* airlines (later we will describe a numerical example that compares these two cases). This result helps to distinguish our analysis from the analysis of Parlar (1988) and of [Lippman and McCardle \(1997\)](#). In the classic 'newsvendor game' there is only one type of demand faced by each firm, a Nash equilibrium always exists and increase in inventory by one firm leads to a decrease in inventory by the competitors (a game of substitutes). In our model, the addition of another demand class competing for the same resource not only complicates the analysis but also raises the possibility that there will not be a joint pure strategy at all, and/or that an increase in the booking limit of one airline may lead to an *increase*

in the booking limit of the other airline. However, we show that under many reasonable demand distributions, such counterintuitive behavior does not appear.

Proposition 2. *Suppose that D_{Li} and D_{Hi}^T , $i=1,2$, are TP_2 . Then a pure strategy Nash equilibrium exists. In addition, the best response functions are decreasing (possibly with jumps down): if one player increases the booking limit, it is optimal for the other player to decrease the booking limit.*

Proof: See the Appendix.

It is worth noting that the proof of Proposition 2 does not rely on showing that the game is submodular (this has been shown for newsvendor games, see Lippman and McCardle 1997), for the game under horizontal competition is not submodular. Instead, we utilize Tarsky's fixed point theorem and explore the structural properties of the best response functions. Note also that proposition 2 specifies conditions on D_{Li} and D_{Hi}^T , although realizations of D_{Hi}^T may depend upon all four demand realizations. The following Corollary describes appropriate conditions on these underlying demand distributions.

Corollary 1. *If $(D_{L1}, D_{L2}, D_{H1}, D_{H2})$ are multivariate TP_2 (MTP_2) in their density functions, then the results of Proposition 2 also hold.*

Proof: From Theorem 2.2 in Joe (1997), if $\underline{D} = (D_{L1}, D_{L2}, D_{H1}, D_{H2})$ are MTP_2 , then so are $(f_1(\underline{D}), f_2(\underline{D}), f_3(\underline{D}), f_4(\underline{D}))$ where f_i are increasing functions. Because D_{Hi}^T is an increasing function of $(D_{L1}, D_{L2}, D_{H1}, D_{H2})$, D_{Li} and D_{Hi}^T are TP_2 as well. The results of Proposition 2 follow. ■

As we mentioned in Section 3, there are many useful distributions that satisfy the MTP_2 property. Multivariate normal distributions of three or more random variables are MTP_2 if the inverse of the covariance matrix satisfies certain properties (specifically, it must be an 'M-matrix', see [Karlin and Rinott, 1980](#)). On the other hand, if a multivariate normal distribution is MTP_2 , then a *necessary* condition is that its correlation coefficients are non-negative. When the demand distributions for high and low-fare passengers (D_{Li} and D_{Hi}) are strongly negatively correlated, then D_{Li} and D_{Hi}^T are not TP_2 , and the Proposition 2 suggests that we may not see a well-behaved game. However, when demand distributions between airlines (e.g., D_{Li} and D_{Lj}) are negatively associated, D_{Li} and D_{Hi}^T may still satisfy the TP_2 condition. In that case, even though $\underline{D}$ is not TP_2 , Proposition 2 may still hold.

Figures 2 and 3, for example, show the best response functions $r_i(B_i)$ of two airlines with $C_i = C_j = 200$ and demands that are characterized by large, negative correlation. To generate Figure 2, we created a demand distribution with a correlation of -0.9 between high and low-fare demands within each airline,

but with independent demands between airlines.² For Figure 3, high and low-fare demands are independent, but the between-airline correlation is -0.9. As suggested by the Proposition 2, the results are very different. Figure 2 shows a game without any pure-strategy equilibrium, and contains a ‘jump’ in which a rise in the booking limit of one airline leads to a sharp rise in the booking limit of the other airline. The game in Figure 3, on the other hand, is well-behaved, with a single equilibrium at a booking limit of 50 seats. In the general horizontal competitive game described above, it is also possible to have more than one equilibrium.

Figure 2: Best response functions with $cor(D_{Li}, D_{Hi}) = cor(D_{Lj}, D_{Hj}) = -0.9$, $cor(D_{Li}, D_{Lj}) = cor(D_{Hi}, D_{Hj}) = 0$

Figure 3: Best response functions with $cor(D_{Li}, D_{Hi}) = cor(D_{Lj}, D_{Hj}) = 0$, $cor(D_{Li}, D_{Lj}) = cor(D_{Hi}, D_{Hj}) = -0.9$

4.2 Comparing the Competitors and a Monopolist We will now compare the behavior of two airlines in competition with the behavior of a monopolist. Note that the term 'monopolist' does not necessarily imply that a single firm is the only airline on a particular route. The 'monopolist' may be two airlines in an alliance to coordinate yield management decisions. In addition, two airlines may compete on a particular route at certain times of day, while each airline may hold a virtual monopoly at other times of day because its competitor has not scheduled a competing flight at a point close in time.

In general, we will find that under horizontal competition the total booking limit for the monopolist is never less than the sum of the booking limits of two competing airlines. To help us understand the dynamics of the game, we first examine a simplified model with high-fare passenger overflow only, and then we discuss numerical experiments with the full model. To enable proper

² Intuitively, jumps in the best responses occur because, under negative correlation, each airline faces two entirely opposite states of the world: (i) heavy low-fare demand and no high-fare demand and (ii) heavy high-fare demand and no low-fare demand. Therefore, each airline establishes either a large booking limit, essentially betting on the first state of the world, or it can establish a low booking limit while hoping for the second state. Under certain problem parameters the two decisions produce the same expected profit so that the objective function is bimodal and the best response is to jump from one mode to the other when the competitor's decision changes (see Appendix, Figure 5). However, such an extreme negative correlation is unlikely to occur in practice. Demand forecast errors for different fare classes are more likely to be independent or weakly positively correlated (see Section 4.3). Therefore, the outcome in Figure 3 is more plausible.

comparison, we assume that the prices p_{ki} and the distribution of consumer demands D_{ki} are equal under the competitive and monopoly environments.

Consider the situation when low-fare customers do *not* overflow to a second-choice flight while high-fare passengers do overflow. This applies in the (unlikely) case that low-fare customers are willing to fly on only one airline. A similar, even simpler model applies to the case where low-fare demand is sufficiently large so that the booking limits of both airlines are *always* reached (e.g., it is high enough to fill an entire aircraft) and therefore low-fare overflow is irrelevant. With no low-fare overflow, objective function (4.1) is replaced by

$$\pi_i = E \left[p_{Li} \min(D_{Li}, B_i) + p_{Hi} \min(D_{Hi}^T, R_i) \right], \quad (4.3)$$

and the optimality condition $\partial \pi_i / \partial B_i = 0$ at B_i^c can be written as,

$$\Pr(D_{Hi}^T > C_i - B_i^c \mid D_{Li} > B_i^c) = \frac{p_{Li}}{p_{Hi}}. \quad (4.4)$$

It can be shown that there is a unique equilibrium in this game (see Netessine and Shumsky, 2004, Proposition S1 for details). In the following proposition we will compare booking limits under competition with system-optimal booking limits. Note that in this proposition and in all subsequent propositions we assume that the optimal solutions for the monopolist and the equilibrium under competition are in the interior, i.e., $B_i^a, B_i^c \in (0, C_i)$.

Proposition 3. *Assume that each airline maximizes objective function (4.3), so that there is no low-fare overflow, and that D_{Hi}^T and D_{Li} are TP₂. Then the booking limits are lower under competition than the centralized solution: $B_i^c \leq B_i^a$ and $B_j^c \leq B_j^a$.*

Proof: The objective function of the alliance is the sum of the two airlines' objective functions, $\pi = \pi_i + \pi_j$, and the centralized optimality condition $\partial(\pi_i + \pi_j) / \partial B_i = 0$ can be written as,

$$\Pr(D_{Hi}^T > C_i - B_i^a \mid D_{Li} > B_i^a) = \frac{p_{Li}}{p_{Hi}} + \frac{\partial \pi_j}{\partial B_i} \bigg|_{(B_i^a, B_j^a)} \frac{1}{p_{Hi} \Pr(D_{Li} > B_i^a)}. \quad (4.5)$$

Clearly, $\partial \pi_j / \partial B_i > 0$ because an increase in the booking limit by one airline results in more high-fare passengers for the other airline without any effect on demand by low-fare passengers. By comparing (4.4) and (4.5) we find,

$$\Pr(D_{Hi}^T > C_i - B_i^c \mid D_{Li} > B_i^c) < \Pr(D_{Hi}^T > C_i - B_i^a \mid D_{Li} > B_i^a). \quad (4.6)$$

Now consider the following four cases:

1. $B_i^c > B_i^a, B_j^c > B_j^a$. Given the TP_2 assumption, the probability term in (4.6) is increasing in both B_i and B_j , so this case is impossible.
2. $B_i^c > B_i^a, B_j^c < B_j^a$. Consider the change from B_i^c to B_i^a . A decrease in B_i leads to a decrease in the probability term while an increase in B_j leads to an increase in the probability term. For inequality (4.6) to hold, the second effect must dominate. But for this to happen we need $|\Delta B_i| < |\Delta B_j|$ where Δ refers to an absolute change in the booking limit. However, analysis of the second optimality condition $\partial \pi_j / \partial B_j$ leads to the opposite requirement $|\Delta B_i| > |\Delta B_j|$ which is a contradiction.
3. $B_i^c < B_i^a, B_j^c > B_j^a$. This leads to another contradiction, by the same reasoning as the previous case.

The only remaining option is that $B_i^c \geq B_i^a$ and $B_j^c \geq B_j^a$. ■

Proposition 3 implies that, under competition, at least as many seats are held for high-fare customers as is optimal under joint profit maximization. For the monopolist, every high-fare passenger who does not find a seat at airline i and turns to airline j is not 'lost' to the firm. Under competition, however, when airline i establishes a lower booking limit, airline j lowers its booking limit as well as the two airlines compete for high-fare passengers.

A similar proof demonstrates that when there is only *low-fare* overflow, and no high-fare overflow, then $B_i^c \geq B_i^a$ and $B_j^c \geq B_j^a$. We do not believe that this result has practical significance, for it is difficult to imagine a situation in which high-fare customers do not overflow (e.g., if there are few high-fare customers, then the booking limit is adjusted accordingly so the overflow is likely to occur). However, this result, when combined with Proposition 3, shows that the outcome of the full game described in Section 4.1 is not clear, for the full game has both types of overflow.

4.3 Numerical analysis To determine whether the previous section's results apply to the full-fledged game, we calculate numerically both the competitive equilibrium and the optimal monopoly solution in a game with both low-fare and high-fare overflow. We perform these numerical experiments over a range of parameter values that is sufficiently wide so that we replicate most possible real-world scenarios. Our goal is to see whether the booking limit set by the monopoly, $B_1^a + B_2^a$, is consistently greater than or

equal to the total booking limit under competition, $B_1^c + B_2^c$.

To find the appropriate range of parameters we examined a variety of sources, including published papers, unpublished PhD theses, and databases collected by the airline industry and the U.S. Department of Transportation (DOT). Our most significant primary source is the DOT's quarterly 'Passenger Origin and Destination Survey' (DOT, 2004). This database, known as either 'Data Base 1a' or 'the O&D Survey,' contains a 10% sample of all airline tickets sold for flights originating during a 3-month period. Each record of each ticket includes the cities visited by the passenger (the itinerary), the price of the ticket, and a *fare code* indicating whether the ticket was restricted or unrestricted. In general, unrestricted, or full-fare, tickets are not subject to limitations such as advance purchase requirements, minimum/maximum stays, or refund penalties. To estimate parameters of our model that are related to fare classes, such as the ratio p_H / p_L , we used the database's fare code as a proxy for our high and low-fare categories. Given that our model allows for only two types of fare classes, we believe that dividing fares into restricted and unrestricted categories is a reasonable approximation.

In total, the O&D Survey provided us with information on 471,000 tickets sold in 1500 different markets (origin/destination pairs). The following list summarizes the data analysis and the parameter values. For more details on the O&D Survey, our analysis of these data, and the following numerical experiments, see Netessine and Shumsky (2004).

- Capacity (C_1 and C_2): Fleet data from the Federal Aviation Administration (FAA, 2004) indicates that the average airplane of a major airline has 180 coach seats, varying from approximately 50 for a regional jet to over 400 for a Boeing 747. Given these data, we design two sets of experiments: a symmetric case with $C_1 = C_2 = 200$, and an asymmetric case with $C_1 = 200$ and $C_2 = 100$.
- Ratio of high fare to low fare (p_H / p_L): In our data set derived from the O&D Survey, the median fare ratio among all 1500 markets is $p_H / p_L = 2.6$. For over 90% of all markets, the ratio fell within a range from 1.3 to 4. Therefore, we define three scenarios, $p_H / p_L = [1.3, 2.6, 4]$, for both the symmetric and asymmetric cases.
- Proportion of demand due to low-fare passengers: Let μ_{Li} (μ_{Hi}) be the average low-fare (high-fare) demand for airline i , $i=1,2$. The proportion of low-fare demand is $\mu_{Li} / (\mu_{Li} + \mu_{Hi})$. In the data from the O&D Survey, the median value among all markets is 0.74, with 90% of all markets falling between 0.5 and 0.9. Therefore, we set $\mu_{Li} / (\mu_{Li} + \mu_{Hi}) = [0.5, 0.74, 0.9]$ for both the symmetric and asymmetric cases.

- Total demand and demand faced by each airline: According to the airline industry trade group the Air Transport Association (ATA), the industry load factor (the utilization of airplane seats) has hovered near 70% for the last decade (ATA, 2004). However, there is substantial flight-to-flight variation around this range, and the application of revenue management techniques generates load factors that are lower than exogenous demand. Here we assume that total demand is equal to total airline capacity. In our experiments, this led to realized load factors between 70% and 98%, depending upon the values of the other parameters.

We must also allocate demand between the airlines. It is most convenient to describe this allocation in terms of the ‘load’ placed on each airline, where the load for airline $i = (\mu_{Li} + \mu_{Hi}) / C_i$. For the symmetric case, we vary airline 1’s load through three parameters, [0.5, 0.75, 1], which implies that airline 2’s complementary load is [1.5, 1.25, 1]. For the asymmetric case, airline 1’s load = [0.5, 1, 1.25] and airline 2’s load=[2, 1, 0.5] (recall that in the asymmetric case, airline 1 has 200 seats and airline 2 has 100).

- Variability: To limit the number of parameters, we assume that all four customer demand distributions have the same coefficient of variation, CV . Based on data analyses in Belobaba (1987) and in Jacobs, Ratliff and Smith (2000), reasonable values for the $CV = [0.2, 0.33, 0.6]$.
- Correlation: Given the data analysis described in Belobaba (1987) and our own discussions with managers who work with yield management systems in the airline industry, correlation in the forecast error among customer classes is usually small. When correlation is significant, positive correlation is probably more prevalent than negative correlation. Therefore, we assume that correlation $\rho = [-0.3, 0.0, 0.3, 0.6]$. To limit the number of parameters, we assume that the correlations among all demands are equal.
- Probability density: Belobaba (1987) uses industry studies and his own data analysis to show that the Normal distribution is a reasonable model of demand within each fare class. For each of our scenarios, we assume that demand is distributed according to a multivariate Normal distribution and is truncated at zero; any negative demand is added to a mass point at zero.

When combined, these parameters define $2 * 3 * 3 * 3 * 3 * 4 = 648$ scenarios. Solutions were found by a simple gradient algorithm and the gradients themselves were evaluated by Monte Carlo integration.

For every scenario, we found that the booking limit set by the monopoly, $B_1^a + B_2^a$, was greater than or equal to the total booking limit under competition, $B_1^c + B_2^c$. The mean difference $(B_1^a + B_2^a) - (B_1^c + B_2^c)$ across all scenarios is 9.3 seats, and the difference varies from 0 seats to 111 seats.

In general, the largest differences occur when correlation is low and when both capacity and demand are balanced among airlines and classes. For a more detailed description of these results, see Netessine and Shumsky (2004).

If we define the ‘service level’ of a customer class as the probability that all customers of a particular class are able to purchase a seat on either aircraft (similar to the ‘no-stockout probability’ of inventory theory), these differences in booking limits can produce significantly different service levels. Over all 648 scenarios, the service level for low-fare customers rose an average of 5.4 percentage points under the monopoly (43% to 48%), while the service level for high-fare customers declined an average of 5.3 percentage points under the monopoly (76% to 71%). The range of results was extremely large. The difference in low-fare service levels was as high as 62%, while the difference in high-fare service levels reached 30%.

To summarize, in this section we demonstrated that the horizontal game may not be well-behaved: best responses may exhibit jumps and, quite counter-intuitively, it may be optimal for an airline to increase its booking limit in response to the increase of booking limit by the competitor. We provided regularity conditions that ensure existence of a competitive equilibrium and guarantee that the best response functions are decreasing. Through a combination of analytical results and numerical experiments, we found that competing airlines tend to reserve too many seats for high-fare passengers. In the conclusion, Section 8, we will discuss additional managerial insights from this section’s analysis, as well as from the following analysis of vertical competition.

5. Vertical Competition

Now consider two or more airlines that do not compete horizontally because all of the airlines operate flights in non-overlapping markets: if airline 1 operates a flight from A to B, then no other airline operates flights between those cities. However, the airlines do interact because they may exchange connecting passengers. For example, if airline 2 flies from B to C, then passengers may travel from A to C by flying on both airline 1 and airline 2. The presence of both *local* passengers flying one leg (from A to B or from B to C), and *connecting* passengers flying both legs leads to important questions of coordination between the airlines. How many high-fare seats should each airline set aside for connecting passengers? For local passengers? What happens if the airlines do not collaborate in these decisions?

Our analysis of the DOT’s O&D Survey found many examples of flights by the major airlines with a high proportion of inter-airline transfer passengers. For example, on flights operated by Continental between Detroit and Cleveland, we estimate that 65% of passengers transferred to or from other airlines. Most of these passengers flew on Northwest on the previous leg and transferred from

Northwest to Continental at Northwest's hub in Detroit. In addition, the data show that regional airlines often carry an even higher proportion of passengers who transfer to or from other airlines.

	low-fare passengers	high-fare passengers	all passengers
<i>Major U.S. Airlines</i>			
American	0.18	0.18	0.18
Continental	0.21	0.21	0.21
Delta	0.13	0.17	0.13
Northwest	0.17	0.21	0.18
United	0.17	0.16	0.17
USAirways	0.16	0.19	0.17
<i>Independent regional airlines</i>			
Alaska	0.35	0.23	0.34
Hawaiian	0.29	0.05	0.25
Horizon	0.45	0.30	0.44
Mesa	0.84	0.64	0.80
Skyking	0.61	0.40	0.59
Skywest	0.93	0.90	0.92
<i>Regional airlines controlled by major airlines</i>			
American Eagle	0.75	0.53	0.69
Comair (Delta)	0.73	0.65	0.71
Delta Connection – Atlanta Southeast	0.84	0.73	0.82
Mesaba/Northwest Airlink	0.84	0.59	0.80
United Express – Atlantic	0.93	0.89	0.92

**Table 1: Fraction of U.S. passengers who are connecting to or from another airline
(source: DOT's O&D Survey)**

Table 1 shows the estimated fraction of inter-airline connecting passengers on flights in the U.S. for many of the largest airlines listed in the DOT's O&D Survey data from the 4th quarter of 1999 (DOT 2004). To calculate these numbers we first examined each flight leg operated by each airline and found the number of passengers whose ticket showed either a transfer from a different airline on the previous leg or a transfer to a different airline on the next leg. The actual numbers in Table 1 were produced by dividing the number of these inter-airline connecting passengers by the total number of passengers. In addition, the O&D Survey classifies each ticket as 'restricted' or 'unrestricted', and we used this classification as a proxy for our 'low-fare' and 'high-fare' categories (for more on this classification scheme, see Section 4.3).

From the Table we see that approximately 15-20% of the passengers on the major airlines are connecting passengers. Regional airlines controlled by the major airlines such as American Eagle and Comair (which is a wholly owned subsidiary of Delta) tend to have a high proportion of connections to and from the controlling partner, around 70% - 90%. For these partners, revenue management decisions are coordinated so that there is no 'vertical competition' between the major airline and the subsidiary. On

the other hand, consider the independent regional airlines, which also have extremely high proportions of connecting passengers but do not coordinate pricing and seat allocation decisions with their partners. These airlines provide us with examples of vertical competition. For example, Skywest feeds passengers to flights operated by Continental, Delta and United, and over 90% of its passengers are connecting passengers. Geographically concentrated airlines, such as Alaska Airlines and Horizon Air also saw a large number of transfers from other airlines, 34% and 44%, respectively. The analysis that led to Table 1 included only flights operated within the U.S. Most international flights operated by the major airlines also have a high fraction of passengers connecting with other airlines (Fernandez de la Torre, 1999).

In this section we consider the impact of these connecting passengers on yield management policies. Throughout this Section we assume that none of the airlines compete horizontally, so that each operates legs in separate markets. Second, we assume that the total revenue generated by any itinerary is the sum of the revenues that would have been generated by a local passenger on each leg of the route. For example, suppose airline 1 flies on a route from city A to city B, while airline 2 flies from B to C. Local passengers traveling from A to B generate revenue p_{k1} , $k = L, H$ to airline 1 and are worth p_{k2} to airline 2. Connecting passengers traveling from A to C are worth $p_{k1} + p_{k2}$, $k = L, H$.

We must also make assumptions about how the revenue is split between the airlines. In our model, the airlines' code sharing agreements allow each airline to book a connecting passenger on any other airline, and if an airline makes a connecting booking it keeps the local fare and distributes the remainder to the appropriate partners. Note that the actual 'local fares' may be negotiated among the airlines in advance. For example, they may be based on the length of the two legs, a system called 'mileage proration' (Boyd, 1998a). However, we assume that the revenue contributed by a connecting passenger to each airline is equal to the revenue that each airline would obtain from a local customer flying a single leg in the same fare class. This simplifies the problem because connecting and local demands then fall naturally into the same booking-limit classes (or buckets) on each airline.

Given this partition of capacity, each airline uses a 'free sale' mechanism to allocate capacity to the other airline. Specifically, each airline provides real-time seat availability information to its partner. If an airline wants to book a low-fare seat on the other airline, it can, as long as the appropriate booking class is open. This mechanism is sometimes called 'automated codesharing,' and it is an increasingly popular mechanism for allocating capacity among airline alliance partners (Fernandez de la Torre, 1999 and Talluri and van Ryzin, 2004).

First, we wish to determine whether the airlines, each optimizing this objective function over its own booking limits, reach a competitive equilibrium. In addition, we want to compare the behavior of these uncoordinated airlines with the behavior of a single, centralized airline that optimizes over the entire

network. In fact, there are a variety of methods that a central authority may use to allocate seats among the various types of demand on each flight leg (recall that each leg has low and high-fare local demands as well as low and high-fare connecting demands along any number of routes). Talluri and van Ryzin (2004) provide a lucid introduction to this topic. They describe (i) 'partitioned booking limits' so that each type of demand has its own booking class, (ii) 'virtual nesting controls' in which flights with similar profit implication are placed in the same virtual class, and (iii) bid-price controls in which each booking request is evaluated in real-time against the total expected marginal cost of satisfying that request.

We will assume that our centralized airline uses a virtual nesting control system and that local and all connecting passengers of each fare class continue to co-exist in the same virtual class. In other words, the centralized airline maintains the same two-class booking limit structure described above; the only difference is that the airline determines a different, 'optimal' booking limit for each flight leg that may take into account the larger revenue that connecting customers can provide. For many parameters this is a reasonable method for allocating capacity. For example, the expected net marginal network profit generated by a low-fare connecting customer is often closer to the local low-fare revenue than the local high-fare revenue (although the connecting passenger who travels on more than one leg generates more revenue, that passenger may displace a local customer on each leg). Such virtual nesting systems are quite popular in practice because of their simplicity and their conformance with existing airline reservation systems. Maintaining the same booking-limit structure also allows us to compare directly the competitive equilibrium and centralized cases.

To gain some intuition about what happens under competition, we will now consider a few extreme cases. In Section 5.1 we examine a general network topology with any number of airlines making any number of connections, but we limit the allocation of connecting customers to either the high-fare or low-fare categories, but not both.³ We will see that in the first scenario competing airlines set booking limits that are lower than is centrally optimal and in the second scenario competing airlines set booking limits that are higher than centrally optimal. In Section 5.2 we construct a model of two airlines exchanging passengers across two legs, and we examine numerically the impact of having passenger connections in both fare categories.

5.1. General networks with connecting customers in either the low-fare or the high-fare bucket

³ One might also consider another special case with both types of connecting demand but no local demand. While this may appear interesting, it turns out that local demand plays an important stabilizing role in the game and that an analysis of this special case is uninformative. For example, consider any number of airlines that operate a sequence of legs connecting cities A to B to C.... It can be shown that when there is no local demand, *any* solution $B_i = B_j, \forall i, j$ such that $0 \leq B_i \leq \min_{vj} (C_j)$, is a Nash equilibrium.

Consider an M -leg network between an arbitrary number of cities. The network may be a series of cities, A to B to C..., or it may be a more complex network involving multiple airlines feeding passengers to multiple hubs. Each leg on this network is controlled by one of N airlines, $N \leq M$. Legs are indexed by $i=1, \dots, M$ and airlines are indexed by $k=1, \dots, N$. Capacity offered on each leg is C_i . Passengers travel on individual legs as well as on any combinations of successive legs. Passengers traveling on a single leg i pay p_{ki} , $k = L, H$ and passengers traveling on a route that includes a set Ω of several legs pay $\sum_{j \in \Omega} p_{kj}$, $k = L, H$. The airlines' code-sharing agreement allows each airline to book a connecting passenger on the other airline, and if an airline controlling legs in Ω makes a connecting booking in class k , it keeps $\sum_{j \in \Omega} p_{kj}$ and pays $\sum_{j \in \Omega/\Omega} p_{kj}$ to the airlines controlling other legs. To describe the total demand faced by each airline, we define $\bar{B}$ as a vector of booking limits on all M legs and $\bar{B}_{-i}$ as a vector of booking limits on $M-1$ legs excluding leg i . Furthermore,

$$D_{Li}^T(\bar{B}_{-i}) = D_{Li} + \bar{D}_{Li}(\bar{B}_{-i}) \text{ and}$$

$$D_{Hi}^T(\bar{B}) = D_{Hi} + \bar{D}_{Hi}(\bar{B}), \text{ where}$$

D_{Li}, D_{Hi} = local low-fare and high-fare demand specific to the leg i ,

$\bar{D}_{Li}(\bar{B}_{-i}), \bar{D}_{Hi}(\bar{B})$ = connecting low-fare and high-fare demand on the leg i , given the decisions made on the other legs.

We assume that the connecting demands have the following properties: $\bar{D}_{Li}(\bar{B}_{-i})$ is increasing in B_{-i} , and $\bar{D}_{Hi}(\bar{B})$ is decreasing in $\bar{B}$. For the analytical results in this section we will not specify a functional form for $\bar{D}_{Li}(B_j)$ or $\bar{D}_{Hi}(B_i, B_j)$, for these first-order assumptions are sufficient. These assumptions are also quite intuitive: a higher booking limit on any one leg opens more seats for connecting low-fare passengers for other airlines on the network, while a lower booking limit opens more seats for connecting high-fare passengers. In Section 5.2 we will describe one specific functional form for the connecting demands that satisfies these assumptions.

The profit generated for an airline from a particular leg i is,

$$\pi_i = E \left[p_{Li} \min \left(D_{Li} + \bar{D}_{Li}(\bar{B}_{-i}), B_i \right) + p_{Hi} \min \left(D_{Hi} + \bar{D}_{Hi}(\bar{B}), C_i - \min \left(D_{Li} + \bar{D}_{Li}(\bar{B}_{-i}), B_i \right) \right) \right]$$

The total profit of airline k controlling a set of legs Ω_k is $\Pi_k = \sum_{j \in \Omega_k} \pi_j$. We will compare the booking limits resulting from such vertical competition with the booking limits that a monopolistic airline would set. For a single monopolistic airline set Ω simply includes all M legs.

We will consider two extreme situations: in the first, all connecting passengers will fly in the low fare only and in the second all connecting passengers will fly in the high fare only. Before we do so, we will prove one theorem that will be instrumental later on.

Proposition 4. *Consider an N -players non-cooperative game with each player $k=1, \dots, N$ endowed with a vector of strategies $\bar{B}_k = (B_{k1}, B_{k2}, \dots, B_{ki}, \dots)$ and a payoff $\Pi_k(\bar{B})$. Let $\bar{B}^c$ denote equilibria of this game and let $\bar{B}^a$ denote the “alliance” solutions. Then, if each $\Pi_k(\bar{B})$ is supermodular in $\bar{B}$ and moreover, if $\partial \Pi_k / \partial B_{ji} \geq 0, j \neq k, \forall i$ (correspondingly, $\partial \Pi_k / \partial B_{ji} < 0, j \neq k, \forall i$) then $\bar{B}^c \leq \bar{B}^a$ ($\bar{B}^c \geq \bar{B}^a$).*

Proof: Consider a fictitious game in which every player is still endowed with a vector of strategies $\bar{B}_k = (B_{k1}, B_{k2}, \dots, B_{ki}, \dots)$ but the payoff now is

$$\bar{\Pi}_k(\bar{B}, \alpha) = \Pi_k(\bar{B}) + \alpha \sum_{m \neq k} \Pi_m(\bar{B}), \forall k. \quad (5.1)$$

Note that for $\alpha = 0$ the equilibrium of this game is $\bar{B}^c$ and for $\alpha = 1$ the equilibrium of this game is $\bar{B}^a$ because each player optimizes the sum of all players' payoffs. Hence, in order to compare the competitive solution with the alliance solution it suffices to show that the equilibrium of the game defined by (5.1) is monotone in α . To show that this is the case, observe that $\bar{\Pi}_k(\bar{B}, \alpha)$ is supermodular in $\bar{B}$ as a sum of supermodular functions. Further,

$$\begin{aligned} \frac{\partial \bar{\Pi}_k(\bar{B}, \alpha)}{\partial \alpha} &= \sum_{m \neq k} \Pi_m(\bar{B}), \forall k, \\ \frac{\partial^2 \bar{\Pi}_k(\bar{B}, \alpha)}{\partial \alpha \partial B_{ji}} &= \sum_{m \neq k} \frac{\partial \Pi_m(\bar{B})}{\partial B_{ji}}, \forall k, j, i. \end{aligned}$$

Hence, if $\partial \Pi_k / \partial B_{ji} \geq 0, j \neq k, \forall i$, we have that $\partial^2 \bar{\Pi}_k(\bar{B}, \alpha) / \partial \alpha \partial B_{ji} \geq 0$ so that the game is supermodular in all players' strategies as well as in the parameter α . From Topkis (1998), it follows that the set of equilibria of this game is monotonically increasing in α and hence $\bar{B}^c \leq \bar{B}^a$. The second part is shown analogously. ■

At this point, it is worthwhile pointing out that $\bar{B}^c \leq \bar{B}^a$ in a set-ordering sense. That is, the largest and

the smallest equilibria satisfy this inequality but not necessarily every solution does. Topkis (1998, pg. 32) calls this the ‘induced set ordering.’ Of course, if the equilibria are unique then $\bar{B}^c \leq \bar{B}^a$ in the conventional way.

We now return to the airline problem. First, consider the problem with connecting passengers in the low fare only. In this case revenues on a given leg are

$$\pi_i(\bar{B}) = E \left[p_{Li} \min(D_{Li} + \bar{D}_{Li}(\bar{B}_{-i}), B_i) + p_{Hi} \min(D_{Hi}, C_i - \min(D_{Li} + \bar{D}_{Li}(\bar{B}_{-i}), B_i)) \right].$$

Proposition 5. *Suppose the distribution of (D_{Li}, D_{Hi}) is TP_2 and all connecting passengers are in the low fare. Then $\bar{B}^c \leq \bar{B}^a$.*

Proof: Note first that all $\pi_i(\bar{B})$ are clearly increasing in $\bar{B}_{-i}$ because an increase in $\bar{B}_{-i}$ makes more connecting low-fare demand available for the leg i without affecting high-fare demand. Further, the first derivative is

$$\begin{aligned} \frac{\partial \pi_i}{\partial B_i} &= p_{Li} \Pr(D_{Li} + \bar{D}_{Li}(\bar{B}_{-i}) > B_i) - p_{Hi} \Pr(D_{Hi} > C_i - B_i, D_{Li} + \bar{D}_{Li}(\bar{B}_{-i}) > B_i) \\ &= \Pr(D_{Li} + \bar{D}_{Li}(\bar{B}_{-i}) > B_i) \left(p_{Li} - p_{Hi} \min(D_{Hi} > C_i - B_i \mid D_{Li} + \bar{D}_{Li}(\bar{B}_{-i}) > B_i) \right). \end{aligned}$$

Under the TP_2 assumption, $\partial \pi_i / \partial B_i$ is increasing in $\bar{B}_{-i}$ so that every $\pi_i(\bar{B})$ is supermodular, and so is $\Pi_k(\bar{B})$. Hence, conditions of Proposition 4 are satisfied and $\bar{B}^c \leq \bar{B}^a$. ■

We now turn to the case with all connecting passengers in the high fare. Recall that, generally speaking, $\bar{D}_{Hi}(\bar{B})$ is a function of B_i as well as $\bar{B}_{-i}$. The dependence on B_i is due to the fact that B_i affects demand for low-fare connecting passengers on legs other than i which, in its turn, affects the number of connecting high fare passengers that can be accommodated which affects $\bar{D}_{Hi}(\bar{B})$. However, if connecting passengers are only in the high fare, the dependence of $\bar{D}_{Hi}(\bar{B})$ on B_i disappears and we can write $\bar{D}_{Hi}(\bar{B}_{-i})$. Revenues on a given leg are

$$\pi_i = E \left[p_{Li} \min(D_{Li}, B_i) + p_{Hi} \min(D_{Hi} + \bar{D}_{Hi}(\bar{B}_{-i}), C_i - \min(D_{Li}, B_i)) \right]$$

Proposition 6. *Suppose that all connecting passengers are in the high fare. Then $\bar{B}^c \geq \bar{B}^a$.*

Proof: Note first that all $\pi_i(\bar{B})$ are decreasing in $\bar{B}_{-i}$ because an increase in $\bar{B}_{-i}$ makes less connecting

high-fare demand available for the leg i without affecting low-fare demand. Further, the first derivative is

$$\frac{\partial \pi_i}{\partial B_i} = p_{Li} \Pr(D_{Li} > B_i) - p_{Hi} \Pr(D_{Hi} + \bar{D}_{Hi}(\bar{B}_{-i}) > C_i - B_i, D_{Li} > B_i).$$

Therefore, $\partial \pi_i / \partial B_i$ is increasing in $\bar{B}_{-i}$ so that every $\pi_i(\bar{B})$ is supermodular, and so is $\Pi_k(\bar{B})$. Hence, conditions of Proposition 4 are satisfied and $\bar{B}^c \geq \bar{B}^a$. ■

The results of Propositions 5 and 6 are intuitive: competing airlines undervalue connecting customers because each airline only receives a fraction of the total revenue from those customers. When connecting traffic travels only in the low-fare bucket, competing airlines establish lower booking limits and sell fewer connecting tickets than is optimal for the network. The opposite happens when connecting passengers travel only in the high-fare bucket. These results also suggest that when connecting passengers are allocated to both buckets, the comparison of booking limits under competition and centralization may depend upon the proportion of connecting passengers in the two buckets, the conjecture that we verify in the next section.

5.2. *Connecting passengers in both buckets.*

In order to verify our conjecture in the previous section, we focus on a simpler network, with just two airlines (airline 1 and airline 2) operating two legs (legs 1 and 2, respectively). We also assign a particular functional form to $\bar{D}_{Li}(B_j)$ and $\bar{D}_{Hi}(B_i, B_j)$. When considering how these demands are generated, the following issue arises: if both local and connecting passengers are accommodated in a bucket, how many seats are sold to local passengers and how many to connecting passengers? For this example, assume that both local and connecting passengers arrive uniformly over the reservation period. First consider low-fare connecting passengers. If the booking limit of airline j is not a constraint, then airline j can book all available connecting passengers, and these passengers are available for airline i : $\bar{D}_{Li}(B_j) = D_L$. However, if the booking limit of airline j is a constraint, the number of local and connecting tickets sold is proportional to their realized demands: $\bar{D}_{Li}(B_j) = (D_L / (D_L + D_{Lj})) B_j$. Similar reasoning applies for high-fare demand. Therefore,

$$\bar{D}_{Li}(B_j) = \min\left(\frac{B_j}{D_L + D_{Lj}}, 1\right) D_L, \quad (5.2)$$

$$\bar{D}_{Hi}(B_i, B_j) = \min \left(\frac{C_i - \min(B_j, D_{Lj} + \bar{D}_{Lj}(B_i))}{D_{Hj} + D_H}, 1 \right) D_H. \quad (5.3)$$

As required, $\bar{D}_{Li}(B_j)$ is increasing in B_j and $\bar{D}_{Hi}(B_i, B_j)$ is decreasing in B_i and B_j . Again, definitions (5.2) and (5.3) are useful examples and will be used in numerical experiments later in this section, but they are not necessary for the analytical results in Section 5.1.

Airline i 's objective function is,

$$\pi_i = E \left[p_{Li} \min(D_{Li} + \bar{D}_{Li}(B_j), B_i) + p_{Hi} \min(D_{Hi} + \bar{D}_{Hi}(B_i, B_j), C_i - \min(D_{Li} + \bar{D}_{Li}(B_j), B_i)) \right]$$

In Table 1, we saw that airlines must often make yield management decisions with a significant number of inter-airline connections in all fare classes. To examine this general case, we conducted numerical experiments to find the difference between the competitive equilibrium booking limits, B_i^c , and the system-optimal booking limits, B_i^a for a variety of problem parameters. Here we will present the results of one set of experiments with two identical airlines, each with 200 seats and each facing a total expected demand of 200 passengers, including both local and connecting passengers in both fare classes. (After the application of yield management controls, this produces load factors between 85% and 95% in the following set of experiment.) We also performed an additional set of experiments with asymmetric airlines: one regional airline with a 100-seat plane ‘feeding’ a major airline with a 200-seat plane. The results of those experiments were similar to the results shown here and hence are not reported.

We again used the O&D Survey and other sources of data to determine the parameters for the experiments. Based on the analysis described in Section 4.3, we set p_H / p_L to the baseline value of 2.6. All demands were distributed as independent normal random variables, truncated at 0, with coefficients of variation equal to 0.33 (Belobaba, 1987). Given the results from the O&D Survey, we allocated 74% of local demand to low-fare passengers. To determine the proportion of high-fare and low-fare connecting traffic, we again examined each of the markets in the O&D survey. We found that among inter-airline connecting passengers, the fraction of high-fare tickets varied from 0 to 0.9. Therefore, we varied the fraction $E(D_H) / E(D_L + D_H)$ across this range. Finally, let f_c represent the fraction of demand due to passengers connecting from other airlines, a number that is roughly equivalent to the statistic shown in the third column of Table 1. We set $f_c = [0.25, 0.75]$, a range that includes most of the markets in the O&D Survey with a significant number of connecting travelers.

In all experiments, we assumed that connecting demands were generated according to the

uniformity assumption that led to equations (5.2) and (5.3). We used Monte Carlo simulation to evaluate the objective function when finding the optimal booking limits for each airline, and variability in the generated demand accounts for the random variation in the trend line. Figure 4 shows the results of the experiments.

Figure 4: Comparing competitive and centralized booking limits

The left-hand side of the figure, with $E(D_H) / E(D_L + D_H) = 0$, is equivalent to the scenario with all connecting passengers in the low-fare bucket (Section 5.1), and we see that equilibrium booking limits are slightly lower under competition which is consistent with our results. When $E(D_H) / E(D_L + D_H) = 0.9$, most connecting passengers are in the high-fare bucket and, as we found in Section 5.1, the booking limits are higher under competition. The intermediate cases show a gradual transition between the two extremes, and at one point the competing airlines establish the centrally optimal booking limits. However, in the majority of examples we examined, the impact of connecting passengers in the high-fare bucket dominated, so that booking limits tend to be higher under competition. From Figure 4 we see that having proportionally more connecting demand (higher f_c) tends to increase the difference between the competitive and centralized booking limits.

In this set of experiments, the difference in network profits between the ‘optimal’ solution and the competitive equilibrium varied from 0, where the solutions are equal, to 0.6% when 75% of all demand is connecting demand, and most of that connecting demand is in the high-fare bin. Higher profit differences were seen with even higher levels of connecting demand and stronger negative correlations among demands. For comparison, Talluri and Van Ryzin (2002) mention that optimal network controls can lead to "improvements on the order of 0.5% ... with gains can be as high as 2% or more under high load

factors." Note, however, that we are not comparing the competitive equilibrium with the truly optimal control scheme; our ‘centralized’ solution is the optimal booking limit, given the sub-optimal virtual nesting scheme that places local and connecting demands in the same fare class.

6. Revenue-sharing Agreements

We have shown that both vertical and horizontal competition on airline routes can lead to a sub-optimal allocation of seats among fare classes. Here we examine how contractual arrangements can coordinate the actions of the two airlines. This is a timely subject, given the proliferation of airline alliances that are presumably trying to do just that. Other work on this subject includes Wynne (1995), who bases transfer prices between airlines on the value of local fares charged in each market traveled by the connecting passengers, Boyd (1998b) who finds prices from the marginal value of each seat on each airline, and Feng and Gallego (2003), who examine contracts in which a firm pays a fixed fare to an airline, while the airline agrees to accept bookings as long as seats are available. In practice, the airline industry has been struggling to develop effective coordination schemes for its various alliances (Fernandez de la [Torre, 1999](#)).

There has been a wealth of research on contracting in the area of supply chain management (see [Cachon, 2004](#), for an overview). Most of these contracts combine a wholesale price with another lever, such as buyback agreements, revenue-sharing agreements, or options. In our setting there is no equivalent of the ‘wholesale price’ between a supplier and a retailer, and therefore we will specify a different type of contract based on transfer prices. However, the method used to derive our contract is similar to methods used to determine many supply chain contracts: we reformulate each player’s objective function so that it becomes a linear transformation of the centralized objective function. This is similar in spirit to the “cost-based revenue sharing” described by Fernandez de la Torre (1999), who notes that such schemes “[raise] the difficult question of how to allocate revenue generated by connecting traffic to the leg that is codeshared.” (pg. 155) That is the problem we address here. Note that this technique is appropriate not only when airlines compete on booking limits, but also when they compete on prices. Here we will derive a suitable contract for airlines under vertical competition (i.e., airlines in a code-sharing alliance). Similar techniques identify contracts that coordinate two airlines under horizontal competition.

Each airline’s objective function is

$$\pi_i = E \left[p_{Li} \min(D_{Li}^T, B_i) + p_{Hi} \min(D_{Hi}^T, C_i - \min(D_{Li}^T, B_i)) \right].$$

where each demand is a function of prices as well as booking limits. During a negotiation phase, the airlines agree to split profits in proportions a and $1-a$. One type of contract found in the supply-chain

literature would simply allocate total revenues to the two airlines according to proportions a and $1-a$. However, such contracts ask the airlines to share revenues from both connecting *and* local traffic, and this may be impossible for technological, competitive, and/or regulatory reasons.

Instead we will assume that the terms of the contract are similar to the scheme proposed by Boyd (1998b) and described in Talluri and van Ryzin (2004, Section 6.2): when an airline sells a connecting ticket, it pays a transfer price to the other airline for the seat. However, our method for finding the transfer payment is different from the duality-based method proposed by Boyd. Here we will demonstrate this technique for sales in the low-fare class; the derivation for the high-fare commission is similar. When airline i sells a connecting ticket, it collects revenue $p_{L1} + p_{L2}$ and pays airline j a transfer price δ_j . The method for calculating the transfer price δ_j is negotiated in advance, and cannot be changed once payments begin. As in the previous section, a local low-fare passenger on airline i is worth p_{Li} in revenue.

In our formulation the transfer price will depend upon the expected flow of traffic, specifically, the expected number of connecting passengers and total passengers, given each airline's booking limit. Let T be the number of connecting tickets sold in the low-fare class and let T_i be the number of connecting tickets sold by airline i (so that $T_1 + T_2 = T$). Given this notation, airline 1 earns $p_{L1}(\min(D_{L1}^T, B_1) - T)$ from local passengers, $(p_{L1} + p_{L2} - \delta_2)T_1$ from connecting tickets sold by airline 1 and $\delta_1(T - T_1)$ from connecting tickets sold by the partner. To find the coordinating contract, we choose δ_1 and δ_2 so that the expected total revenue from low-fare customers for airline 1 is equal to proportion a of the centralized revenue from low-fare customers:

$$\begin{aligned} & E\left[p_{L1}(\min(D_{L1}^T, B_1) - T) + (p_{L1} + p_{L2} - \delta_2)T_1 + \delta_1(T - T_1)\right] \\ &= E\left[a(p_{L1} \min(D_{L1}^T, B_1) + p_{L2} \min(D_{L2}^T, B_2))\right]. \end{aligned} \quad (5.4)$$

Now let β_i be the ratio of the expected number of connecting passengers and the expected number of passengers accommodated by airline i : $\beta_i = E[T_i] / E[\min(D_{Li}^T, B_i)]$. Also let λ be the ratio of the expected number of connecting tickets sold by airline 1 and the total number of connecting tickets sold: $\lambda = E[T_1] / E[T]$. Dividing (5.4) throughout by $E[\min(D_{L1}^T, B_1)]$,

$$p_{L1}(1 - \beta_1) + (p_{L1} + p_{L2} - \delta_2)\lambda\beta_1 + \delta_1(\beta_1 - \lambda\beta_1) = a\left(p_{L1} + p_{L2} \frac{\beta_1}{\beta_2}\right).$$

Rearranging terms and dividing through by β_1 ,

$$(1-\lambda)(\delta_1 - p_{L1}) - \lambda(\delta_2 - p_{L2}) = a \frac{p_{L2}}{\beta_2} - (1-a) \frac{p_{L1}}{\beta_1}. \quad (5.5)$$

Equation (5.5) shows how the transfer prices depend upon the traffic flows. Given a value of $\delta_2 > p_{L2}$, (5.5) indicates that the transfer price δ_1 (i) increases as λ increases (ii) increases as β_1 increases and (iii) decreases as β_2 increases. Point (i) is intuitive: as airline 1 sells a larger proportion of connecting traffic, it receives fewer transfer payments from airline 2, and a larger transfer price is needed to align its behavior. Points (ii) and (iii) are more subtle. For point (ii), a rise in β_1 indicates that connecting tickets represent a higher proportion of airline 1's sales. This implies that airline 1 is paying airline 2 a transfer price on a higher proportion of its own ticket sales because of a decrease in airline 1's local demand. Because airline 1 acts in its own self-interest, a fall in local demand would push airline 1 to *lower* its booking limit at the expense of connecting traffic on airline 2. Therefore, the incentive for airline 1 to preserve room for connecting traffic must rise, and δ_1 rises as β_1 rises. The reason for point (iii) is similar.

7. Observations and Future Research

In this paper we have examined how booking-limit decisions are affected by both horizontal competition (with passenger overflow) and vertical competition (with connecting passengers). We have shown that a simple condition on the demand distribution, total positivity of order 2, is sufficient to ensure a pure strategy Nash equilibrium under horizontal competition and under vertical competition when connecting demand is restricted to the high or low-fare classes. In general, we find that the equilibrium behavior of competing airlines can be very different from the behavior of competitors in a 'newsvendor game.' In this simpler game, each competitor's objective function is concave and the game is submodular. However, in our setting neither property holds, and this leads to best-response functions that can be discontinuous, and can be either decreasing or increasing. The reason for this disparity is that in the newsvendor game the decision is how much inventory/capacity to procure while in our model the decision is to allocate *fixed* capacity among two customer classes. The problem of allocating fixed capacity is not always well-behaved even in the absence of competition, i.e. the objective function may not be unimodal, resulting in the possible lack of competitive equilibrium when competition is introduced. Our model demonstrates that computerized revenue management systems may generate capacity controls that exhibit non-intuitive behavior, e.g., may suggest a significant increase in a booking limit without any significant change in demand, although such behavior is unlikely. An understanding of the underlying causes of these effects

can be useful for managers.

We have also found that booking limits are lower under horizontal competition than the booking limits found by a central profit-maximizer. Competitive equilibrium booking limits may be higher, or lower, under vertical competition, although in a 2-airline, 2-leg case we found that they tend to be higher. The differences in booking limits can lead to significant differences in the level of service among customer classes. These results can be useful for managers who are planning expansion into new markets or facing an entry by a rival. Although airlines lose revenue because of competition, our results indicate that some groups of customers (high-fare customers under horizontal competition and low-fare customers under vertical competition) gain because their service is higher under competition than under a monopoly. The net value to consumer welfare of these losses and gains are, however, not clear, for it depends upon the comparative value those customers place on the ability to obtain tickets in their chosen fare classes.

We have demonstrated analytically and evaluated numerically both the direction and magnitude of revenue losses for the two airlines due to competition. Effective coordinating contracts can reduce these losses, and while there has been much research activity in the area of supply chain contracting, the same has not occurred in the yield management area. Here we have described a set of revenue-sharing contracts to coordinate the booking-limit decisions of two airlines with connecting passengers. Our work is an initial attempt in this direction. Given the proliferation of airline alliances and prevalence of competition on airline routes, additional efforts are needed.

Our models of full competition included only two airlines with two passenger classes each. In practice, multiple airlines can be in competition and typically more than two fare classes are offered by each. Relaxing each of these assumptions presents certain challenges. For example, a model with more than two airlines competing horizontally is hard to analyze because one has to specify the order in which passengers overflow when one airline runs out of seats (e.g., from airline 1 to airline 2 and then to airline 3 or first to airline 3 and then airline 2). Consideration of more than two passenger classes complicates the analysis as well because the number of decision variables (booking limits) and hence the number of optimality conditions would grow accordingly. Because direct analysis is difficult, simulation has been used to examine the impact of competition ([Belobaba and Wilson, 1997](#)). On the other hand, in Section 5.1 we saw that the analysis of vertical competition on multiple legs among multiple airlines does not pose significant difficulties, as long as connecting passengers belong to only one fare class. However, when there are multiple fare classes for connecting passengers, the simple supermodular structure of the game is not preserved. In this case, the impact of competition depends upon the proportion of customers in each class. Williamson (1992) and Fernandez de la Torre (1999, pg. 185) use simulation to examine the impact of various optimization techniques and coordination mechanisms in airline networks with a

larger number of fare classes.

Another significant concern with the analysis is that when comparing competitive and cooperative booking limits we assume that both prices and exogenous demand are constant. For some comparisons this assumption is reasonable. Under horizontal competition, two competing airlines often charge the same prices throughout the day for travel on a particular route, and some hours in the day are 'competitive' while others are monopolized by a single airline. Prices are uniform over all flights, but the timing of flights throughout the day affects the yield management decisions of both airlines. Likewise, in a vertical alliance, prices may be negotiated far in advance, while the yield management decisions happen in real-time, according to rules that may be quite similar to those described here.

Finally, in practice airlines may be competing on the same leg (horizontally) in addition to competing on different legs (vertically), and many alliances involve both horizontal and vertical coordination. In our data analysis we did identify examples of horizontal competition with relatively little vertical competition (e.g., shuttle flights between Boston and New York) and vertical competition without significant horizontal competition (e.g., a regional airlines with a monopoly in a small market that is feeding the hub of a major airline). However, the analysis of simultaneous horizontal and vertical competition is an interesting area for additional research.

Appendix: Proof of Proposition 2

To prove existence, we will employ Tarsky's fixed point theorem which states that sufficient conditions for the existence of a pure strategy equilibrium are that the strategy space is closed and bounded and that the best-response functions are non-decreasing ([Vives, 2000](#)). Note that Tarsky's theorem does not require that the best responses are continuous: discontinuities (or 'jumps') are allowed, but only jumps *up*.

Evidently, the strategy space $[0, C_i] \times [0, C_i]$ is closed and bounded. Instead of showing that best responses are non-decreasing, we will show that the best-response functions are non-increasing because after a simple re-definition, $\bar{B} = -B_j$, Tarsky's result applies to this new game ([Vives, 2000](#), page 33, Remark 13). Consequently, jumps down in the best-response functions are allowed, but not jumps up.

We will first use the Implicit Function Theorem (IFT) to show that best responses are decreasing whenever they are differentiable (Part I). This, however, does not eliminate the existence of jumps up because the IFT can only characterize the best response at points where it is differentiable. In Part II we will demonstrate that the best-response functions can only have jumps down.

Part I. By the IFT,

$$\frac{\partial B_i}{\partial B_j} = - \frac{\partial^2 \pi}{\partial B_i \partial B_j} / \frac{\partial^2 \pi}{\partial B_i^2} \quad (5.6)$$

At player i 's best response, $\partial^2 \pi_i / \partial B_i^2 < 0$ and we will show that $\partial^2 \pi_i / \partial B_i \partial B_j < 0$ whenever the first-order conditions hold. Equivalently, we want to show that the first derivative (4.2) is decreasing in B_j . First note that, given a particular realization $\underline{D} = (D_{Li}, D_{Lj}, D_{Hi}, D_{Hj})$, D_{Li}^T is decreasing in B_j and D_{Hi}^T is increasing in B_j . The first two probability terms can be rewritten as,

$$\begin{aligned} & p_{Li} \Pr(D_{Li}^T > B_i) - p_{Hi} \Pr(D_{Hi}^T > C_i - B_i, D_{Li}^T > B_i) \\ &= \Pr(D_{Li}^T > B_i) \left[p_{Li} - p_{Hi} \Pr(D_{Hi}^T > C_i - B_i \mid D_{Li} > B_i - (D_{Lj} - B_j)^+) \right]. \end{aligned} \quad (5.7)$$

Because D_{Li}^T is decreasing in B_j , the first probability term $\Pr(D_{Li}^T > B_i)$ is decreasing in B_j . From Theorem 2.3 of Joe (1997), TP₂ implies that D_{Li} and D_{Hi}^T are right-tail increasing, i.e., $\Pr(D_{Hi}^T > C_i - B_i \mid D_{Li} > B_i - (D_{Lj} - B_j)^+)$ is increasing in B_j . Finally, since the first-order conditions hold at the equilibrium, we are assured that the expression in square brackets is positive and therefore (5.7) is decreasing in B_j .

Now we show that the last probability term in (4.2) is increasing in B_j and, because this term is multiplied by $-p_H$, the derivative must be decreasing in B_j . The final term can be rewritten,

$$\begin{aligned} & \Pr(D_{Li} > B_i, D_{Lj}^T < B_j, D_{Hj} > R_j, D_{Hi}^T < C_i - B_i) \\ &= \Pr(D_{Li} > B_i, D_{Li} + D_{Lj} < B_i + B_j, D_{Li} + D_{Lj} + D_{Hj} > C_j + B_i, D_{Li} + D_{Lj} + D_{Hi} + D_{Hj} < C_i + C_j) \end{aligned}$$

The booking limit B_j only appears once in this expression, in the event ' $D_{Li} + D_{Lj} < B_i + B_j$ ', and the probability of this event increases as B_j increases. Therefore, the last probability term increases, $\partial^2 \pi_i / \partial B_i \partial B_j < 0$, and best responses are decreasing whenever they are differentiable.

Corollary to Part I (will be used in Part II): Suppose $B_i^*(B_j^*)$ is some point on player's i best response function. Then $\partial^2 \pi_i / \partial B_i \partial B_j < 0$ holds for any $(B_i, B_j) \in [0, B_i^*] \times [0, B_j^*]$.

Proof: We only need to show that

$$p_{Li} > p_{Hi} \Pr(D_{Hi}^T > C_i - B_i \mid D_{Li} > B_i - (D_{Lj} - B_j)^+) \quad (5.8)$$

everywhere on $(B_i, B_j) \in [0, B_i^*] \times [0, B_j^*]$, and then the reasoning of Part I applies to show that

$\partial \pi_i / \partial B_i$ is decreasing in B_j over the appropriate range. We demonstrated that (5.8) is true at $B_i^*(B_j^*)$.

We know that the probability term on the right is increasing in B_j if TP_2 holds so the result holds for any $B_j \leq B_j^*$. Similarly, we can show that this term is increasing in B_i if TP_2 holds. Expand this term as,

$$\Pr \left(B_i + D_{Hi} + \left(D_{Hj} - C_j + \min \left(D_{Lj} + (D_{Li} - B_i)^+, B_j \right) \right)^+ > C_i \middle| D_{Li} > B_i - (D_{Lj} - B_j)^+ \right)$$

There are three B_i in this expression. An increase in the first or third values of B_i clearly leads to an increase in the probability expression, but the impact of the second B_i is the opposite. However, the first B_i dominates the second; the expression

$$B_i + D_{Hi} + \left(D_{Hj} - C_j + \min \left(D_{Lj} + (D_{Li} - B_i)^+, B_j \right) \right)^+$$

is increasing in B_i . Therefore, if the inequality (5.8) holds at $B_i^*(B_j^*)$, it will also hold at any $B_i < B_i^*$ and $B_j < B_j^*$. $\square$

Part II: We now eliminate the possibility of jumps *up* in the best responses. Jumps occur when the objective function is bi-modal (or multi-modal). We will only consider the bi-modal case (the multi-modal case is analogous). The proof is by contradiction. Suppose that there is a jump up in the best response $B_i(B_j)$ that occurs at B_j^* (see Figure 5). There are two global maxima $B_i^1(B_j^*) < B_i^2(B_j^*)$, $\pi_i(B_i^1, B_j^*) = \pi_i(B_i^2, B_j^*)$, and the first-order necessary condition holds at each of them. Because there is a jump up, $\pi_i(B_i^1, B_j^* + \varepsilon) < \pi_i(B_i^2, B_j^* + \varepsilon)$. That is, as we increase B_j^* infinitesimally, the second maximum becomes the unique global maximum. This implies,

$$d\pi_i/dB_j \big|_{(B_i^1, B_j^*)} < d\pi_i/dB_j \big|_{(B_i^2, B_j^*)},$$

because the objective function of player i rises faster (or declines more slowly) at the second local maximum B_i^2 than at the first local maximum B_i^1 . Using the envelope theorem,

$$\partial \pi_i / \partial B_j \big|_{(B_i^1, B_j^*)} < \partial \pi_i / \partial B_j \big|_{(B_i^2, B_j^*)}. \quad (5.9)$$

A necessary condition for (5.9) is that $\partial^2 \pi_i / \partial B_i \partial B_j > 0$ somewhere between the two best response points

Figure 5: A Jump Up

$B_i^1(B_j^*)$ and $B_i^2(B_j^*)$. However, from the Corollary to Part I we know that $\partial^2 \pi_i / \partial x_i \partial x_j \leq 0$ for every $(B_i, B_j) \in [0, B_i^2(B_j^*)] \times [0, B_j^*]$, and this is sufficient to show a contradiction. Therefore, a jump up is impossible, and the best-response function is decreasing. $\square$

Acknowledgements

We would like to thank Andy Boyd, Zhihang Chi, Gerard Cachon, Marshall Freimer, Martin Lariviere, Phillip Lederer, Bill Lovejoy, Michael Raith, two anonymous referees and an associate editor for helpful comments. We also thank Laura X.L. Liu for her help with the data analysis.

References

- Air Transport Association (ATA). 2004. Annual Traffic and Capacity – U.S. Airlines Scheduled Service. Available at <http://www.airlines.org/econ/d.aspx?nid=1032>.
- Belobaba, P.P. 1987. Air travel demand and airline seat inventory management, Ph.D. Thesis, Massachusetts Institute of Technology.
- Belobaba, P. 1989. Application of a probabilistic decision model to airline seat inventory control. *Operations Research*, vol. 37, no. 2, 183-197.
- Belobaba, P. 1998. The Evolution of Airline Yield Management: Fare Class to Origin Destination Seat Inventory Control. In *Handbook of Airline Marketing*, Butler, G. F. and M. R. Keller, eds., McGraw-Hill. pp. 285-302.
- Belobaba, P. and J. L. Wilson. 1997. Impacts of yield management in competitive airline markets. *Journal of Air Transportation Management*, vol. 3, no. 1, 3-9.
- Borenstein, S. and J. Netz. 1999. Why do all the flights leave at 8 am?: Competition and departure-time differentiation in airline markets. *International Journal of Industrial Organization*, vol. 17, 611-640.
- Borenstein, S. and N. L. Rose. 1994. Competition and price dispersion in the US airline industry. *Journal of Political Economy*, vol. 102, no. 4, 653-683.

- Boyd, A. E. 1998a. Airline Alliances. *OR/MS Today*, October 1998.
- Boyd, A. E. 1998b. Airline Alliance Revenue Management. PROS Strategic Solutions Technical Report, PROS Revenue Management, 3100 Main Street, Suite #900, Houston, TX 77002.
- Brueckner J. K. and W. T. Whalen. 2000. The price effects of international airline alliances. *Journal of Law & Economics*, Vol. XLIII, 503-545.
- Brueckner, J. K. 2003. International airfares in the age of alliances: the effects of codesharing and antitrust immunity. *Review of Economics and Statistics*, vol. 85, 105-118.
- Brumelle, S. L., J. I. McGill, T. H. Oum, K. Sawaki and M. W. Tretheway. 1990. Allocation of airline seats between stochastically dependent demands. *Transportation Science*, vol. 24, no. 3, 183-192.
- Cachon, G. 2004. Supply Chain Coordination with Contract. In the *Handbook of Operations Management*, edited by Steve Graves and Ton de Kok.
- Cachon, G. and M. Lariviere. 2000. Supply chain coordination with revenue-sharing contracts: strengths and limitations. Forthcoming in *Management Science*.
- Department of Transportation (DOT). 2004. Sources of Aviation Data. U. S. Department of Transportation, The Office of Airline Information. Available at http://www.bts.gov/programs/airline_information/sources/.
- Dobson G. and P. J. Lederer. 1993. Airline scheduling and routing in a hub-and-spoke system. *Transportation Science*, vol. 27, no. 3, 281-97.
- Federal Aviation Administration (FAA). 2004. FAA Flight Standards Service, Aviation Information Website. Available at <http://av-info.faa.gov/OpCert.asp>.
- Feng, Y. and G. Gallego. 2003. Evaluating Long-Term Contracts for Airlines. Working paper, Columbia University.
- Fernandez de la Torre, P. E. 1999. Airline Alliances: The Airline Perspective. Masters Thesis, Massachusetts Institute of Technology. Available at theses.mit.edu/Dienst/UI/2.0/Composite/0018.mit.theses%2f1999-291/1?nsections=9
- Hotelling, H. 1929. Stability in competition. *Economic Journal*, vol. 39, 41-57.
- Joe, H. 1997. Multivariate Models and Dependence Concepts. Chapman & Hall, London.
- Karlin, S. and Rinott, Y. 1980. Classes of orderings of measures and related correlation inequalities I. Multivariate totally positive distributions. *J. Multivariate Anal.*, 10, 467-498.
- Lederer P. J. and Nambimadom R. S. 1998. Airline network design. *Operations Research*, vol. 46, no. 6, 785-804.
- Li, M. Z. F. and T. H. Oum. 1998. Seat allocation game on flights with two fares. Working Paper, Nanyang Technological University.
- Lippman, S. A. and K. F. McCardle. 1997. The Competitive Newsboy. *Operations Research*, vol. 45, no. 1, 54-65.
- Mahajan, S. and G. J. van Ryzin. 2001. Inventory competition under dynamic consumer choice. *Operations Research*, Vol. 49, No. 5, 464-657.
- McGill, J. I. and G. J. van Ryzin. 1999. Revenue management: research overview and prospects. *Transportation Science*, vol. 33, no. 2, 233-256.
- Morrison, S. A. and C. Winston. 1995. *The Evolution of the Airline Industry*. The Brookings Institution, Washington, D.C.
- Netessine, S. and N. Rudi. 2003. Centralized and competitive inventory models with demand

- substitution. *Operations Research*, Vol. 51, No.2, 329-335.
- Netessine, S. and R. Shumsky. 2004. Revenue management games: vertical and horizontal competition. On-line Companion, University of Rochester. Available at <http://mansci.pubs.informs.org/ecompanion.html>.
- Parlar, M. 1988. Game Theoretic Analysis of the Substitutable Product Inventory Problem with Random Demands. *Naval Research Logistics*. Vol. 35, 397-409.
- Richard, O. 2003. Flight Frequency and Mergers in Airline Markets. *International Journal of Industrial Organization*, Vol. 21, 907-922.
- Rudi, N. 2000. Some models of risk pooling. PhD dissertation, University of Pennsylvania.
- Spengler. 1950. Vertical Integration and Anti-Trust Policy. *Journal of Political Economy*. Vol. 58, 347-352.
- Talluri K. and G. van Ryzin. 2004. The theory and practice of revenue management. Kluwer Academic Publishers.
- Tong, Y. L. 1990. *The Multivariate Normal Distribution*. Springer-Verlag, New York.
- Topkis, D. M. 1998. *Supermodularity and Complementarity*. Princeton University Press, Princeton, New Jersey.
- U.S. Department of Transportation. 2000. Multinational Airline Alliances, in *Airline Mergers and Alliances*, Organisation for Economic Cooperation and Development, DAFE/CLP(2000)1, available at www.oecd.org.
- Vives, X. 2000. Oligopoly Pricing: Old Ideas and New Tools. The MIT Press, Cambridge, Massachusetts.
- Williamson, E. L. 1992. *Airline Network Seat Inventory Control: Methodologies and Revenue Impacts*. PhD thesis, Flight Transportation Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts.
- Wynne, R. 1995. From Queen Airs to RJs: evolution in major-regional airline alliances and code sharing. In *Handbook of Airline Economics*, Aviation Week Group.
- Zhao, X. and D. Atkins. 2002. Strategic Revenue Management under Price and Seat Inventory Competition. Working Paper, University of British Columbia.
- Zuckerman, L. 2002. Antitrust Laws an Issue In Airline Alliances. *New York Times*; New York, N.Y.; Apr 28, 2002, pg. 5.3.