

Decoupling Airline Alliances: A Bid-Price Approach to the Game of Incomplete Information

Christopher P. Wright
Simon Graduate School of Business,
University of Rochester, Rochester, New York 14627
wrightc1@simon.rochester.edu

May 31, 2010

The revenue management problem for airlines utilizing interline itineraries, such as those commonly used within alliances, is considerably more complex than that faced by individual airlines managing isolated networks. Additionally, while alliances seek to coordinate the partners' actions, the partners within the alliance may be unwilling or unable to share complete information about their operations because of competitive, technical and legal barriers. Both of these factors suggest that such agreements will usually represent a game of incomplete information.

Wright et al (2010) formulate a model for a two-partner alliance in which one partner can sell an interline itinerary only after paying a posted transfer price to its partner for the relevant subitinerary. They derive the equilibrium decision rules for the static proration scheme prevalent in practice, as well as for several dynamic schemes based on suggestions from practitioners. Their formulation assumes a game of complete information: within each time period, each airline knows its partner's inventory level, and both have identical forecasts of future arrival probabilities and revenue distributions over the entire alliance network. The authors leave open the research question of how such policies could be managed when the partners are faced with incomplete information at the time they make decisions. We begin to address that problem here.

For dynamic transfer price (proration) schemes to function, each partner must post a list of current transfer prices for each interline itinerary. In this paper, we assume an extreme case: that these transfer prices are the only information that each airline has about the state of its partner's network. Specifically, its partner's forecasts (revenue and arrival processes) and inventory levels are not shared. Therefore, each partner's decisions are based solely on the transfer prices, their own forecasts for their own networks, and their own inventory levels.


Furthermore, we assume that the partners utilize the *bid price scheme* first suggested by Vinod (2005) and later analyzed by Wright et al (2010). This scheme is being considered by practitioners because it utilizes bid prices, a component of many modern revenue management systems. However, it also provides structural advantages that we exploit in our method.

We demonstrate that under this new game with incomplete information, the alliance can be decoupled into single airline problems connected by the exchanged bid prices. We examine the decision rules and equilibrium policies under this game and compare the policies, as well as the overall performance of the alliance, to both the centralized control (first-best) alliance and an alliance that shares complete information. We show that it can perform close to or even above the complete information bid price scheme. Its performance does, however, suffer in those cases when the complete information bid price scheme has been shown to be weakest.

Key words: revenue management, yield management, airlines, Markov game, alliances, incomplete information

1. Introduction

A passenger traveling from Rochester, NY in the United States to Sidney, Australia might be surprised to find that she can purchase the multi-leg international itinerary from a single airline. However, several airlines appear to offer exactly that. For example, American Airlines offers the following itinerary:



What appears to be a single-carrier itinerary is, in fact, operated by three different airlines. The first leg – from Rochester, NY to Chicago, IL – is operated by the regional carrier American Eagle. The second leg – from Chicago, IL to Los Angeles, CA – is operated by American Airlines, itself. The final leg – from Los Angeles, CA to Sidney, Australia – is operated by another international carrier, Qantas.

All three airlines are members of the **oneworld** alliance, a global alliance of airlines, but their relationships differ quite substantially. The relationship between American Eagle and American Airlines, both wholly-owned subsidiaries of AMR Corporation, is relatively simple and not the focus of this paper. However, since both American and Qantas are large international carriers, their relationship is significantly more complex.

In Wright et al (2010), the authors lay out a framework for the latter type of relationship, providing a Markov game model, analyzing several revenue management schemes and pointing out some important considerations in scheme design. An assumption of their work is that airlines have complete information about their partners' networks, including structure, inventory levels and forecasts.

While alliances seek to coordinate the partners' actions, the partners may be unwilling or unable to share complete information about their operations because of competitive, technical and legal barriers. For example, they may not provide complete information about demand forecasts

to each other because this would require sharing information on all flights and itineraries that they operate, including those itineraries on which they compete.

Even if the partners are willing and able to share this information, the computational load for such problems would be unmanageable. It is generally accepted that the exact form of even a single airline's problem is beyond the computational capabilities of available technology. For example, the state space for a single airline with 10 flights, each with 100 seats, would have 101^{10} , or just over 10^{20} , possible states in each period. For more realistic airline sizes, the number of states can be substantially larger. As such, much of the revenue management literature has focused on approximations for this problem. Even a simple alliance between two airlines exacerbates this problem. Two airlines of the size above would face 10^{40} states, in addition to a greatly increased number of potential itineraries.

These two factors bring to question how airlines within an alliance can still achieve a high level of revenue without complete information sharing. In this paper, we provide a novel approach that can be utilized by airlines, which requires sharing limited information and allows them to reduce the computational load through the use of existing single-network approximation methods. Our preliminary numerical studies suggest that the scheme could perform well, often on par with the complete information case, but also show that its performance closely follows that of the bid price scheme on which it is based, even when that scheme performs poorly.

The remainder of the paper is organized as follows: We begin with a brief summary of the relevant literature in §2. In §3, we summarize the model described in Wright et al (2010), before describing the limited information version of the game and our heuristic solution in §4. Then, we numerically simulate a selection of sample alliances, showing the potential of the scheme to work well and comparing the effects of utilizing some existing single-network approximation methods. In §5, we briefly describe an extension to our basic method with both potential benefits and drawbacks. Finally, in §6, we provide conclusions and suggestions for future research.

2. Literature Review

From its earliest days, starting with the seminal work of Littlewood (1972), the basic revenue management problem for single airlines has received considerable attention in the literature. A

survey paper by McGill and van Ryzin (1999) and later a book by Talluri and van Ryzin (2004) describe the extensive body of work to date.

However, while several researchers – examples include Wynn (1995), Boyd (1998a), Fernandez de la Torre (1999), Netessine and Shumsky (2005), Vinod (2005), and Shumsky (2006) – have looked at the alliance revenue management problem, it was not until recently that a detailed model of the problem was provided and analyzed. Wright et al (2010) provide the general model of a two-airline alliance, formulating the revenue management problem as a Markov game. They present and analyze several dynamic sharing schemes and compare their performance to the static scheme commonly found in practice. They assume a game of complete information and leave open the research question of how such schemes may be managed under a game of incomplete information. We begin to address that question in this paper.

Game theory, in general, within the revenue management problem has been relatively limited. Vulcano et al (2002) examine a game between an airline and its customers in an auction setting. Netessine and Shumsky (2005) and Raza and Akgunduz (2008) look at games between two competing airlines. Boyd (1998b) describes conditions under which an alliance can be coordinated under a static linear program formulation, utilizing a softblocking scheme.

Games of incomplete information (see Fudenberg and Tirole 1991 for a detailed explanation), specifically, in revenue management have also been relatively rare. Isler and Imhof (2008) look at the single airline revenue management and pricing decisions of competing airlines under incomplete information, while Song (2007) focuses on the hotel revenue problem.

3. Summary of General Alliance Network Model

Before introducing our limited-information game, we must briefly review the dynamic model of a two-partner alliance introduced in Wright et al (2010). The alliance modeled consists of two partner airlines, loosely denoted c and $-c$. While larger alliances consisting of more partners are certainly in wide-spread use, two airlines is sufficient to model most interline itineraries and we will show that the decomposition allows for unlimited extension.

The alliance network is composed of m flights (origin, destination, and departure time) with the number of flights operated by airline c denoted m_c . The m -dimensional vector of remaining inventory for the alliance, $\bar{x} \equiv \{x^1, \dots, x^{m_1}, x^{m_1+1} \dots x^m\}$, can be segregated into individual airline

$J_k^c(\bar{x})$ denotes the total (current and future) expected value for airline c given inventory \bar{x} with k remaining periods. As such, airline c 's opportunity cost of the inventory used by an itinerary is,

$$\Delta J_k^{cj}(\bar{x}) \equiv J_k^c(\bar{x}) - J_k^c(\bar{x} - A^j),$$

We define $J_k^c(\bar{x}) = -\infty$ whenever a component of \bar{x} is negative and, thus, the opportunity cost when \bar{x} is not sufficient to sell an itinerary is $\Delta J_k^{cj}(\bar{x}) = \infty$.

In order for its partner, airline $-c$, to sell an interline itinerary, it must pay airline c the current posted transfer price, $p_k^{cj}(\bar{x})$. Let $\bar{p}_k^c(\bar{x})$ be the n -vector of all transfer prices in period k .

Finally, in each period, airline c determines a set of acceptance rules, $u_k^c(\cdot; \bar{x})$, such that $u_k^{cj}(r; \bar{x}) = 1$ indicates that airline c is willing to accept a request for itinerary j with associated net revenue r – i.e., less any applicable transfer price paid to a partner – in period k with remaining inventory \bar{x} .

Thus, the general form of the Bellman equations can be written as,

$$J_k^c(\bar{x} | u_k^{-c}, \bar{p}_k^{-c}(\bar{x})) = \left(\begin{array}{l} \sum_{j \in N_c} q_k^{cj} \mathbb{E} \left[R_k^{cj} u_k^{cj}(R_k^{cj}) + J_{k-1}^c(\bar{x} - A^j u_k^{cj}(R_k^{cj})) \right] + \\ \sum_{j \in N_s} q_k^{cj} \mathbb{E} \left[(R_k^{cj} - p_k^{-cj}(\bar{x})) u_k^{cj}(R_k^{cj} - p_k^{-cj}(\bar{x}), \bar{x}) \right. \\ \left. + J_{k-1}^c(\bar{x} - A^j u_k^{cj}(R_k^{cj} - p_k^{-cj}(\bar{x}), \bar{x})) \right] + \\ \sum_{j \in N_s} q_k^{-cj} \mathbb{E} \left[p_k^{cj}(\bar{x}) u_k^{-cj}(R_k^{-cj} - p_k^{cj}(\bar{x}), \bar{x}) \right. \\ \left. + J_{k-1}^c(\bar{x} - A^j u_k^{-cj}(R_k^{-cj} - p_k^{cj}(\bar{x}), \bar{x})) \right] + \\ \sum_{j \in N_{-c}} q_k^{-cj} \mathbb{E} \left[J_{k-1}^c(\bar{x} - A^j u_k^{-cj}(R_k^{-cj})) \right] + q_k^0 J_{k-1}^1(\bar{x}) \end{array} \right), \quad (3)$$

$$J_0^c(\bar{x}) = 0 \quad \forall \bar{x} \geq \bar{0}$$

$$\text{where } u_k^{cj}(r, \bar{x}) = \arg \max_{u \in \{0,1\}} \{ru + J_{k-1}(\bar{x} - A^j u)\} \quad \text{and} \quad \tilde{R}_k^{cj}(\bar{x}) = R_k^{cj} - p_k^{-cj}(\bar{x}).$$

with a corresponding set of equations for airline $-c$. Since the choice of c is arbitrary, we will focus only on airline c 's equations.

The first summation in (3) corresponds to airline c 's intraline itinerary requests, the second and third summations correspond to airline c 's and airline $-c$'s interline itinerary requests, respectively, and the final summation corresponds to airline $-c$'s intraline requests. The final term cor-

responds to the “no arrival” case. In each case, one airline – the *marketing airline* – receives a request and decides to accept or not based on the corresponding revenue and, for interline itineraries, the transfer price that it has to pay to its partner – the *operating airline* – for the required legs on its network.

4. A Bid Price Only Game

In this section, we describe an extreme form of the game of incomplete information between the alliance partners, in which information shared is limited to the current posted transfer-prices. We, then, consider a specific form of transfer-pricing scheme, which has structural advantages for solving this problem.

4.1. Basic Game

In order for dynamic transfer price (proration) schemes to function, each partner must post a list of current transfer prices for each interline itinerary. In this paper, we assume an extreme case: that these transfer prices are the only information that each airline has about the state of its partner’s network. Specifically, its partner’s forecasts (revenue and arrival processes) and inventory levels are not shared. Therefore, each partner’s decisions are based solely on the posted transfer prices, its own forecasts for its own network, and its own inventory level.

Furthermore, we will assume a specific class of transfer-pricing schemes: those that require posting the current bid prices. In particular, we focus on one particular scheme, bid price. We summarize this scheme below, but direct the reader to Wright et al (2010) for a complete discussion.

4.2. Bid Price Scheme

Under the *bid price* scheme, each airline posts its true opportunity cost (the bid price) for each interline itinerary in the alliance. That is,

$$p_k^{cj}(\bar{x}) = \Delta J_{k-1}^{cj}(\bar{x})$$

In theory, the operating airline is therefore indifferent about whether its partner sells an interline itinerary, because it is exactly compensated for the inventory used.

In Wright et al (2010), the authors showed that such a scheme is *centrally optimal* for interline itineraries, meaning that the optimal acceptance decision is exactly that of a centralized controller making all acceptance decisions for the alliance. However, since no transfer payment is made for intraline itineraries, airlines do not take into account the effect that selling an intraline itinerary has on its partner. As such, those decisions are not centrally optimal and over the horizon, the alliance may not achieve first-best (those achieved by the centralized controller) revenue levels. Additionally, because the operating airline receives no share of the surplus revenue – those revenues exceeding the combined opportunity costs of the airlines – the bid price scheme fails to pass information about future, high-value requests and may perform poorly when there is a large discrepancy between the revenues and/or arrival probabilities of each partner for certain interline itineraries.

The choice of this scheme here, however, was dependent on two critical factors: (1) practitioners are actually considering this scheme because it utilizes an element of their existing revenue management systems, bid prices, and has a logical appeal; and (2) because of the aforementioned indifference, airlines can ignore interline arrivals to their partners, as the value of either decision – accept or deny – is the same.

4.3. Constant Bid Price (CBP) Heuristic

For this new game with incomplete information, the Bid-Price Only game, we generate a heuristic, taking advantage of the aforementioned structure, as well as a critical result from the literature, that can be utilized by airlines to manage revenues effectively despite the minimal level of information sharing. Under this new heuristic, we are able to modify the Bellman equations from the full information case in such a way that the problem is decoupled with the exception of the current posted bid prices. With this minimal connection, airlines can solve the alliance problem without depending on information about their partners' networks. We examine the airlines' equilibrium policies and compare the overall performance of the alliance, to both the centralized control (first-best) alliance and an alliance that shares complete information.

4.3.1. Decoupling the Alliance Under CBP

We begin by demonstrating the sequence of steps that are used to decouple the original problem of complete information into a set of individual problems with minimal information sharing. As mentioned in §4.2, the advantage of the bid price scheme is that airlines can ignore arrivals of interline itinerary request to their partners, because they are indifferent with respect to their partners' acceptance decisions in such cases. Formally, the third term of the Bellman equations shown in (3) can be replaced by the value going forward with no change in inventory, $J_{k-1}^c(\bar{x})$. The corresponding arrival probability can, therefore, be simply added to the final "no arrival" term, resulting in the following modified Bellman equations,

$$J_k^c(\bar{x} | u_k^{-c}, \bar{p}_k^{-c}(\bar{x})) = \left(\begin{aligned} & \sum_{j \in N_c} q_k^{cj} \mathbb{E} \left[R_k^{cj} u_k^{cj}(R_k^{cj}) + J_{k-1}^c(\bar{x} - A^j u_k^{cj}(R_k^{cj})) \right] + \\ & \sum_{j \in N_s} q_k^{cj} \mathbb{E} \left[\begin{aligned} & (R_k^{cj} - p_k^{-cj}(\bar{x})) u_k^{cj}(R_k^{cj} - p_k^{-cj}(\bar{x}), \bar{x}) \\ & + J_{k-1}^c(\bar{x} - A^j u_k^{cj}(R_k^{cj} - p_k^{-cj}(\bar{x}), \bar{x})) \end{aligned} \right] + \\ & \sum_{j \in N_{-c}} q_k^{-cj} \mathbb{E} \left[J_{k-1}^c(\bar{x} - A^j u_k^{-cj}(R_k^{-cj})) \right] + \left(q_k^0 + \sum_{j \in N_s} q_k^{-cj} \right) J_{k-1}^1(\bar{x}) \end{aligned} \right), \quad (4)$$

Under the assumption of the bid price only game, partners do not have specific information about each other's inventory levels or forecasts. As such, we decompose the inventory vector \bar{x} into those legs operated by airline c , \bar{x}_c , and those operated by its partner, \bar{x}_{-c} . Since airline c only knows its own inventory level, any decision made by its partner for intraline requests will have no effect on the known state of the system for airline c . These decisions will, however, have an indirect effect by changing the posted bid prices.

Using \hat{J} to indicate that this value function is found under limited information, the Bellman equations are modified as such,

$$\hat{J}_k^c(\bar{x}_c | \bar{p}_k^{-c}) = \left(\begin{aligned} & \sum_{j \in N_c} q_k^{cj} \mathbb{E} \left[R_k^{cj} u_k^{cj}(R_k^{cj}, \bar{x}_c) + \hat{J}_{k-1}^c(\bar{x}_c - A^j u_k^{cj}(R_k^{cj}, \bar{x}_c)) \right] + \\ & \sum_{j \in N_s} q_k^{cj} \mathbb{E} \left[\begin{aligned} & (R_k^{cj} - p_k^{-cj}) u_k^{cj}(R_k^{cj} - p_k^{-cj}, \bar{x}_c) \\ & + \hat{J}_{k-1}^c(\bar{x}_c - A^j u_k^{cj}(R_k^{cj} - p_k^{-cj}, \bar{x}_c)) \end{aligned} \right] + \left(q_k^0 + \sum_{j \in N_{-c} \cup N_s} q_k^{-cj} \right) \hat{J}_{k-1}^1(\bar{x}_c) \end{aligned} \right), \quad (5)$$

Finally, the equations in (5) are linked by a series of dependent bid prices passed between the partner over the remainder of the horizon. Without knowledge of the current inventory or any forecast of future arrivals, the values of the bid prices from its partner are no longer perfectly determinable by airline c . For all $t < k$, the value function must be estimated using the expected value of the bid prices given the information set available at time k ,

$$\hat{J}_t^c(\bar{x}_c | \bar{p}_k^{-c}) = \left(\begin{array}{l} \sum_{j \in N_c} q_t^{cj} \mathbb{E} \left[R_t^{cj} u_t^{cj} (R_t^{cj}, \bar{x}_c) + \hat{J}_{t-1}^c(\bar{x}_c - A^{cj} u_t^{cj} (R_t^{cj}, \bar{x}_c)) \right] + \\ \sum_{j \in N_s} q_t^{cj} \mathbb{E} \left[\left(R_t^{cj} - E[p_t^{-cj} | \bar{p}_k^{-c}] \right) u_t^{cj} (R_t^{cj} - E[p_t^{-cj} | \bar{p}_k^{-c}], \bar{x}_c) \right] \\ + \hat{J}_{t-1}^c(\bar{x}_c - A^{cj} u_t^{cj} (R_t^{cj} - E[p_t^{-cj} | \bar{p}_k^{-c}], \bar{x}_c)) \right] \\ + \left(q_k^0 + \sum_{\substack{j \in \\ N_{-c} \cup N_s}} q_t^{-cj} \right) \hat{J}_{t-1}^1(\bar{x}_c) \end{array} \right), \quad (6)$$

How can these bid prices be estimated given the limited information set? While several possible methods could be used, we turn to the literature for our solution. In Akan and Ata (2009), the authors demonstrate that for a variant of the single network problem, the (adapted) bid prices form martingales. Intuitively, this means that airlines manage sales such that the decrease in value due to the passage of time is balanced with the increase in value due to sales. Logically, one would expect that were bid price anticipated to rise or fall, then that information would already be incorporated in the current value.

Given the decoupling of the problem, one would expect that an airlines bid prices adapted to its partner's as they update would then (at least approximately) form a martingale as well. Accepting that airlines cannot easily estimate the effect of changes in their own bid prices on their partners', we take this assumption one step further by assuming that the expected value of any future bid price, conditional on the current posted values, will in fact be those current values. I.e.,

$$E[\bar{p}_t^{-cj} | \bar{p}_k^{-cj}] = \bar{p}_k^{-cj} \quad 1 \leq t \leq k, \quad (7)$$

for all sample paths. Learning, therefore, is limited to the adjustments that their partners make as time passes, inventory changes and their own bid prices change.

With this assumption, the net revenue in any future period follows the same distribution as the gross revenue from a request with the mean shifted by the current posted bid price for that itinerary. For convenience, we introduce a net revenue variable for each remaining period in the horizon conditional on the current posted bid prices, $\bar{R}_t^{cj}(\bar{p}_k^{-c})$, such that,

$$\bar{R}_t^{cj}(\bar{p}_k^{-c}) = \begin{cases} R_t^{cj} - p_k^{-cj} & j \in N_S \\ R_t^{cj} & j \in N_c \end{cases}, \quad (8)$$

Substituting these adjusted revenues into (6) leaves us with a final and familiar form of the Bellman equations for each airline,

$$\hat{J}_t^c(\bar{x}_c | \bar{p}_k^{-c}) = \left(q_t^{c0} \hat{J}_{t-1}^1(\bar{x}_c | \bar{p}_k^{-c}) + \sum_{\substack{j \in \\ N_c \cap N_S}} q_t^{cj} \mathbf{E} \left[\bar{R}_t^{cj}(\bar{p}_k^{-c}) u_t^{cj}(\bar{R}_t^{cj}(\bar{p}_k^{-c}), \bar{x}_c) \right. \right. \\ \left. \left. + \hat{J}_{t-1}^c(\bar{x}_c - A^{cj} u_t^{cj}(\bar{R}_t^{cj}(\bar{p}_k^{-c}), \bar{x}_c)) | \bar{p}_k^{-c} \right] \right), \quad (9)$$

where $q_t^{c0} \equiv q_t^0 + \sum_{\substack{j \in \\ N_c \cap N_S}} q_t^{-cj}$

which is solved from $t = 1$ back to k to generate the current bid prices.

Equation (9) matches the value function for a single airline network (see for example, Talluri and van Ryzin 1998) with revenues equal to the net revenues and the “no arrival” probability equal to q_t^{c0} , the probability that there is not an arrival to airline c in period t .

In our game, at the start of each period, the partners’ revenue management systems iteratively exchange bid prices, hopefully reaching an equilibrium. Acceptance decisions are then made as usual, using the posted equilibrium bid prices.

4.3.2. Additive Bid-Prices

Before continuing with analysis of the method described above, we briefly show that a familiar form of bid prices can be utilized in (9). To this point, we have assumed the most general form of bid price control: a resource *or set of resources* is assigned a value, usually its opportunity cost, and an itinerary is accepted if and only if the revenue exceeds the total assigned value of

the resources required. Formulated as such, the inventory for even a multi-leg itinerary can be assigned a bid price.

The most common form of bid-price controls, however, assigns a value to each resource individually and the opportunity cost of an itinerary is approximated by the sum of the bid prices for the individual legs of which it is composed (see Talluri and van Ryzin 1998 for a detailed explanation.) Formally,

$$p_k^{cj}(\bar{x}) = \Delta J_k^{cj}(\bar{x}) \approx \sum_{i \in A^j} A_{ij} \lambda_k^{ci}(\bar{x}), \quad (10)$$

where $\lambda_k^{ci}(\bar{x})$ is airline c 's the bid price for leg i in period k . The expression in (9) remains unchanged with the exception that the acceptance rule for itinerary j takes on a new form,

$$u_t^{cj}(r; x) = \arg \max_{u \in \{0,1\}} \left\{ \left(r - \sum_{i \in A^{cj}} a_{ij} \lambda_k^{ci}(\bar{x}) \right) u \right\}$$

where $\{i \in A^{cj}\} \equiv \{i \in M_c : a_{ij} \geq 1\}$

Airline c , therefore, needs only track one bid price for each leg on its own network and does not need to know which interline itinerary is being sold, only its required subitinerary. As such, several itineraries may have the same bid price, a potential source of inefficiency in this design.

If the partner, $-c$, utilizes a similar method of bid-price control, then the price paid for its subitinerary will also be a sum of individual bid prices. In such cases, the airlines might choose to simply exchange leg bid prices, summing them up when request a multi-leg subitinerary.

Logically, one might expect that the current opportunity cost of leg i ,

$$\lambda_k^{ci}(\bar{x}) \equiv J_k^c(\bar{x}) - J_k^c(\bar{x} - e_i).$$

would lead to a good approximation of the itinerary's true opportunity cost. However, as the Talluri and van Ryzin (1998) show, these prices are not universally optimal. In fact, one can easily show that they can be very poor estimates of the opportunity cost of inventory, especially at the boundaries of the problem, where opportunity costs can change rapidly with changes in state (inventory and time.)

For example, consider a case where there is one seat left on each of A and B, and a guaranteed arrival in period 1 for itinerary AB at \$1000. Since selling either leg before period 1 would preclude the sale of the AB request, the opportunity costs of each leg for decisions made in period 2 should be \$1000. As such, under the additive bid-price mechanism, a request in period 2 for AB would only be sold if its revenue exceeds the sum of these two bid prices, \$2000. Clearly, though, the opportunity cost of the inventory used is only \$1000.

Despite these limitations, because of the practical convenience and the diminishing effects of these boundary cases on the overall performance as the problem size increases, much of the literature and many systems in use are based upon this assumption.

4.4. Equilibrium Results

Here, we analyze the potential existence and uniqueness of equilibrium posted bid prices in each state of the system. With the assumption that $F_k^{ij}(r)$ is continuous in r , we get the following result:

THEOREM 4.1. For the general alliance, in each period, k , and for each inventory level, \bar{x} , there exists at least one equilibrium set of posted prices.

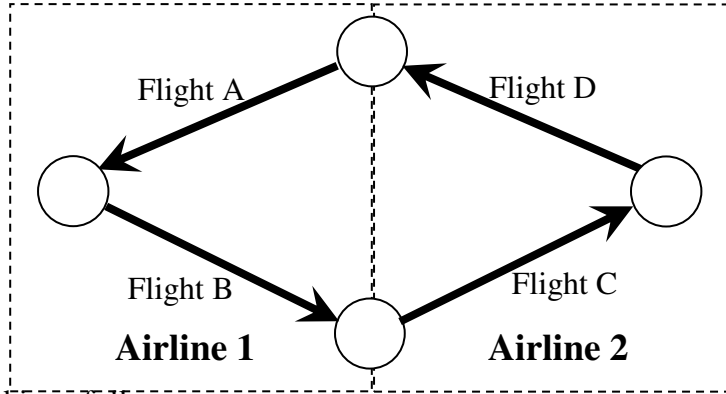
Proof: See Appendix 1.

Uniqueness of the equilibrium bid prices, however, cannot be guaranteed for the general alliance as demonstrated by the counterexample in the following section.

4.4.1. Counterexample to Uniqueness

Consider the alliance shown in Figure 1, in which airline 1 operates flights A and B, while Airline 2 operates flights C and D. Both airlines offer itineraries ABC, BCD, CDA and DAB (all interline itineraries).

Figure 1 – Sample alliance for uniqueness counterexample



In each period, demand is as follows:

Itinerary	Airline 1		Airline 2	
	Prob[Arrival]	Payout	Prob[Arrival]	Payout
ABC	0.01	100	0.24	100
BCD	0.24	100	0.01	100
CDA	0.24	100	0.01	100
DAB	0.01	100	0.24	100

For $k = 3$ with $\bar{x} \equiv \{1, 1, 1, 1\}$, there are two equilibria:

$$\bar{p}_k^1(\bar{x}) \equiv \{\$110.59, \$110.59, \$55.94, \$55.94\}, \bar{p}_k^2(\bar{x}) \equiv \{\$2.59, \$2.59, \$2.59, \$2.59\}$$

and

$$\bar{p}_k^1(\bar{x}) \equiv \{\$2.59, \$2.59, \$2.59, \$2.59\}, \bar{p}_k^2(\bar{x}) \equiv \{\$55.94, \$55.94, \$110.59, \$110.59\}$$

Likewise, in $k = 12$ with $\bar{x} \equiv \{3, 3, 3, 3\}$, there are also two equilibria:

$$\bar{p}_k^1(\bar{x}) \equiv \{\$117.83, \$117.83, \$59.57, \$59.57\}, \bar{p}_k^2(\bar{x}) \equiv \{\$0.0621, \$0.0621, \$0.0621, \$0.0621\}$$

and

$$\bar{p}_k^1(\bar{x}) \equiv \{\$0.0621, \$0.0621, \$0.0621, \$0.0621\}, \bar{p}_k^2(\bar{x}) \equiv \{\$59.57, \$59.57, \$117.83, \$117.83\}$$

In this case, there exist two distinct (opposite) asymmetric equilibria, one high-low and one low-high, in which one partner sells its itineraries with high probability because of the low prices of the other, leading it to have such a high opportunity cost that its partner is unlikely to sell any itinerary, leading to its low prices. However, with only slight changes to the problem parame-

ters, the problem returns to a case with a unique equilibrium, suggesting that this may problem may be limited to unusual and/or boundary cases.

Additionally, a keen reader may have noticed that because the graph of the network is cyclical the counter-example presented above requires time travel in order to accommodate the provided itineraries. While we were unable to establish a proof of uniqueness for a network with an acyclical graph, we were not able to find a counterexample, either.

Note: in the above counterexample, we have used discrete revenue distributions for convenience, but similar results hold for some continuous distributions with sufficiently small standard deviations.

4.4.2. Uniqueness in a Special Case

In the previous section, we showed by counter example that uniqueness of the equilibrium cannot be established for the general network. However, we can however establish uniqueness for at least one type of alliances:

THEOREM 4.2. For an alliance network consisting single leg operated by each carrier with corresponding intraline itinerary and a single interline itinerary combining the two legs, in each period, k , and for each inventory level, \bar{x} , there exists a unique equilibrium set of posted prices.

Proof: See Appendix 1.

Additionally, uniqueness has been demonstrated numerically for several other alliances, an example of which can be seen in the numerical results found in §4.6.

4.5. Benefits and Drawbacks of CBP

Because of its design and the resulting decoupled problem, CBP has several strengths that would make its use advantageous to airlines. It does, however, also have some inherent drawbacks that must be considered when using this scheme.

A primary advantage of the system is one common to all bid-price based schemes: it relies on values, bid prices, that are already a component of many modern revenue management systems.

Even with booking limit control schemes, the acceptance policies often imply a vector of bid prices for the resources. This commonality makes its implementation easier to manage, as the reservation system simply passes along the values that it has already determined.

Likewise, incorporating the partner's posted bid prices should be relatively simple for existing systems to handle. For airlines already utilizing bid prices control systems, in fact, the addition would be trivial. Future revenues used when determining its own bid prices would be adjusted by its partners current bid prices and instead of accepting a request if the revenue exceeds the bid prices for its *own* required legs, the system would accept only if it exceeds the bid prices for its *all* required legs. While less straightforward, for booking limit systems, the adjustment might best be handled by an adjustment to the revenue directly.

Perhaps the greatest advantage of CBP, however, is its reduced computational load versus the original complete information game. It is well accepted that even the single-airline network revenue management problem is computationally prohibitive in its exact form. Without an approximation scheme such as CBP, practitioner would be hopeless to implement any alliance revenue sharing scheme that requires exact value calculations. The savings particular to CBP come in three forms.

First, when determining current bid prices, each airline considers only its own network's resources. The schemes suggested under complete information by Wright et al (2010), theoretically, require the partners to each evaluate its partner's value function for each state of the alliance problem. The inventory portion of the state space for the exact problem is then the product of the capacities for every flight. By contrast, under CBP, the inventory portion of the state space for each airline is only equal to the product of the capacities on its own legs. Thus, for airlines with networks of similar size, the inventory state space – the product of the capacities for every flight – considered by each airline is reduced to roughly the square root of the problem analyzed in the game of complete information.

Second, because additional partners only increase the number of itineraries and revenue adjustments, the inventory state space does not change with the number of partners in the alliances. As daunting as the calculations for a two-airline alliance would be, generally alliances consist of multiple partners, each adding (in product) to the size of the inventory state space. Under CBP, it actually remains a constant size: the product of the capacities for only those flights within an

airline's own network. There would, of course, be some additional load for managing more itineraries and adjustments, but the size of the state space would not grow.

Third, because the resulting decoupled problem exactly matches the single network problem, existing approximation methods can be utilized in calculating these bid prices to further reduce the load. As noted, even the single-airline network revenue management problem is intractable in its exact form. However, a substantial portion of the existing literature is dedicated to estimating the bid prices for a single airline. Such methods include, but are not limited to, the deterministic linear programming (DLP) model, the randomized DLP (RDLP) model, Displacement Adjusted Virtual Nesting, Expected Marginal Seat Revenue (see Talluri and van Ryzin 2004 for more information on those four techniques), dynamic bid price control (Adelman 2007) and Lagrangian decomposition (see Topaloglu 2009, as well as Kunnumkal and Topaloglu 2010).

Combining these advantages certainly increases the likelihood of airlines successfully implementing such a system. However, a few concerns remain that practitioners will need to consider when implementing CBP. As with the strengths, these drawbacks include those inherent to bid-price schemes in general and those specific to CBP.

Perhaps the biggest of these concerns is the "trust" issue. Discussions with practitioners revealed that this concern has been a barrier to implementing any bid-price based system and it will need to be addressed by researchers and practitioners. These schemes are, however, being tested within the largest alliances. The size of both the partners and the alliance itself may make compliance easier to ensure, as the potential loss from having a partner accused of "cheating" the system would be substantial.

As one would expect, CBP seems to follow the performance of the full-information bid price scheme with some loss due to the reduced information sharing. In §4.6, we show a situation in which the bid price scheme performs well and CBP performs similarly well. Our experiments also show that in cases when bid price performs poorly – as in the case illustrated in Figure 4 where bid price does not share information about future interline requests – CBP performs even worse. However, the cases when bid price performs well seem more common than when it does not.

Two additional weaknesses are specific to CBP. First, CBP does not *directly* consider stock-outs of a partner's inventory, as it is assumed that inventory will be available at the current trans-

fer price for any future sales. *Indirectly*, though, as inventory decreases, we would expect the transfer prices to rise, so partners will count on selling less.

Second, while the size of the problem solved has been greatly reduced, it is important to remember that it must be solved iteratively in each period. Of course, one would expect that for larger state spaces (inventory levels and time period remaining), the bid prices would not change substantially, allowing few iteration for convergence. Also, we show an example in §4.6 where bid prices converge very quickly and this rapidness has been observed across numerous example problems.

4.6. Numerical Experiments

In this section, we use small sample alliances to demonstrate the performance of the CBP heuristics. These samples are not sufficiently large to draw major conclusions, but serve to illustrate the potential value of this scheme under various circumstances.

We begin by comparing two small networks similar to those examined with the complete information dynamic schemes analyzed in Wright et al (2010): *bid price*, as described above, uses the current bid prices as the transfer price; *bid-price proration* prorates the revenue by the relative size of the partners' bid prices – i.e., airline c 's share is its bid prices divide by the sum of the two airlines' bid prices; and *partner price* allows both partners to choose any transfer prices they want. We provide the results of those three schemes as a baseline against which we compare the performance of CBP.

In both networks, each carrier operates one flight – A for airline 1 and B for airline 2 – and a single interline itinerary, AB, is offered in addition to the two intraline itineraries. The first network is *symmetric* in that both airlines see the same demand for their intraline flights as their partners do for their own, and the interline demand is split evenly, with the same distribution of arriving revenue to each. The second network is *asymmetric* in that only one airline has intraline demand for its leg and the interline demand is not (necessarily) split evenly.

We begin with the symmetric network. Each intraline has a starting mean revenue, conditional on a request, of \$200, with a standard deviation of \$20. The interline itinerary has a starting mean revenue of \$500 and a standard deviation of \$40. Each of the means grow by 20% over the 30 period horizon and each flight begins with 10 seats. Results are generated for differ-

ent prevalence of intraline versus interline request, ranging from 100% interline to 100% intraline (split evenly between A and B in all cases.)

Figure 2 – CBP vs. Complete-Information (CI) Schemes for Symmetric Alliance

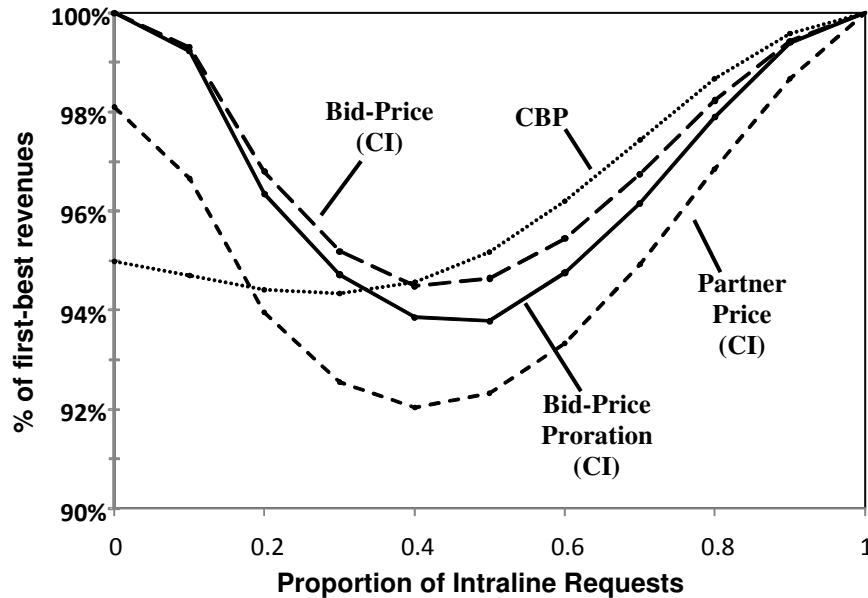


Figure 2 shows, on the right, that when the proportion of intraline requests is high relative to that of intraline itineraries, CBP performs very well relative to the full information case. While it does outperform the complete information schemes in this example, we caution that small changes in the system parameters can shift the ordering. However, for number of such cases, the performances remained similar to the complete information versions. This effect is not entirely surprising, as the amount of information that is not shared, which is already covered in part by indirect effects through the partner’s bid prices, is relatively small in these cases.

On the left, when interline requests dominate the alliance, the effects of not sharing complete information, including not knowing when the partner will stock out or how the bid prices will actually deviate on individual sample paths, drive the performance of CBP down. Below approximately 25%, the difference in performance between CBP and the complete information schemes becomes significant.

Data from the U.S. Department of Transportation (BTS, 2006) shows that, while growing, that the fraction of interline itineraries within the U.S. was at just 20% in 2004. This number suggest that for many cases alliance are likely to the right of this cutoff.

We continue with the asymmetric network. Here, just one intraline itinerary, A, has demand. It has a starting mean revenue of \$200, with a standard deviation of \$20. The interline itinerary has a starting mean revenue of \$400 and a standard deviation of \$40. Again, each of the means grow by 20% over the 30 period horizon and each flight begins with 10 seats. The probability of a request for A and for AB is 50% in each period. Results are generated for different distributions of the interline requests to the two partners, with the percent going to airline 1 ranging from 0% to 100%.

Figure 3 – CBP vs. Complete-Information (CI) Schemes for Asymmetric Alliance

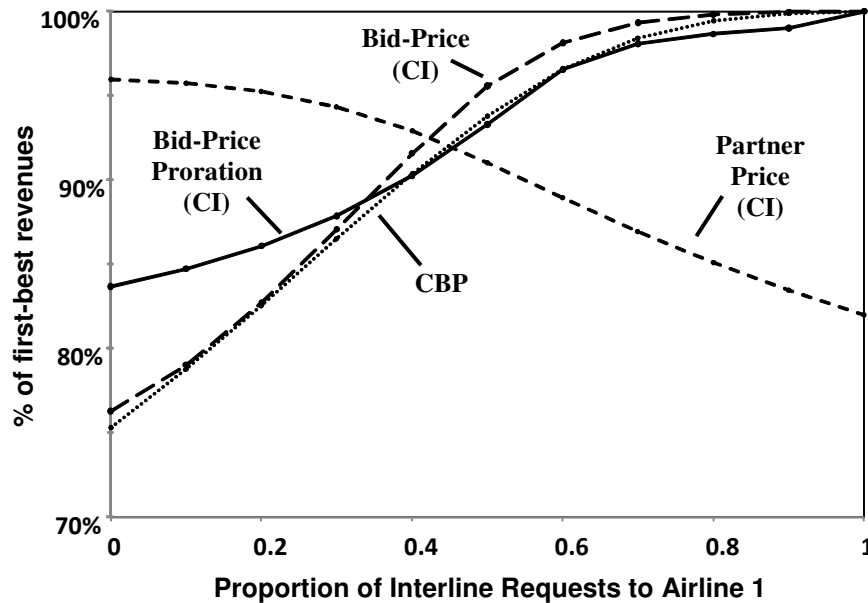
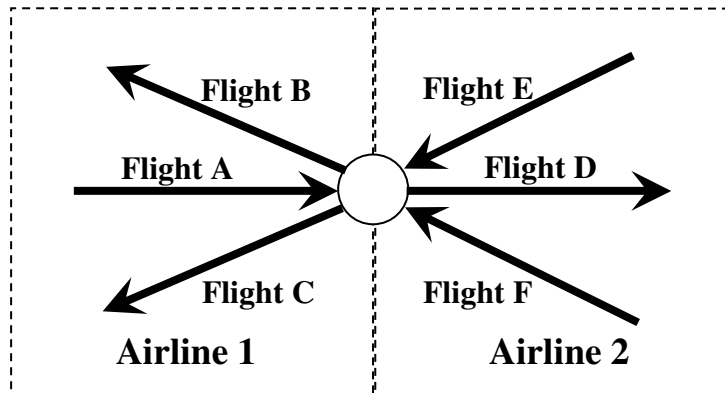


Figure 3 shows, again, that CBP maps closely with the complete information bid-price scheme. However, in this case, that property is “for better or worse.” The left side, where the intraline and interline itineraries are sold by different airlines, airline 2 must provide incentive for airline 1 to hold its A inventory for more valuable AB requests. Wright et al (2010) show that in such cases, schemes that share surplus revenue can provide such incentive. Partner price, or in essence cheating under bid price, actually provides the largest share of the interline revenue

to the operating airline, airline 1, so it performs best in this case. Bid-price proration shares a small amount of the surplus revenue with the operating airline, so its performance is poorer. Bid price and, by extension, CBP share (approximately) none of the surplus revenue and, as such, perform the worst. The relative performance of bid-price proration to bid price in this case is part of the impetus for looking at the Prorated CBP scheme described in §5.

Lastly, consider the alliance shown in Figure 4, in which airline 1 operates flights A, B, and C, while Airline 2 operates flights D, E, F. Each airline sells its own one and two-leg local itineraries, and both airlines offer interline itineraries, AD, EB, EC, FB and FC. While still small, this represents a larger network and we will use to examine some of the effects of increasing the problem size on the performance of CBP.

Figure 4 – Sample alliance for CBP



In each period, a request for each one-leg itinerary arrives with probability 0.04 and for each two-leg itinerary (local and interline) with probability 0.07. Interline requests are split 50/50 between carriers. In the base model, there are seven periods and each flight has one available seat. The scale factor multiplies each of these values. For example, at scale factor 6, the problem consists of 42 periods and each leg has 6 seats.

Figure 5 – Performance of CBP vs. Full-Information for Increasing Problem Size (Scale Factor)

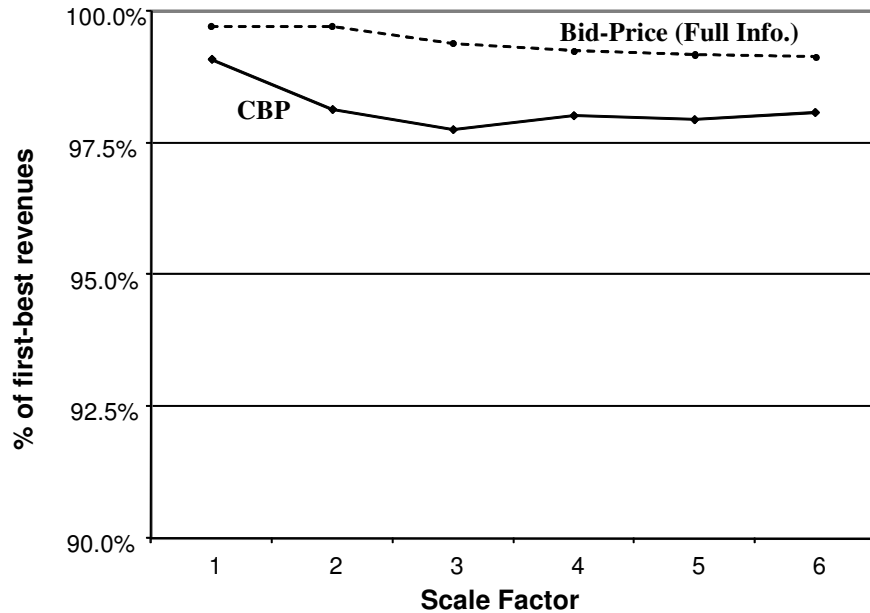



Figure 5 compares the performance of the CBP heuristic to the complete information bid price scheme, each as a percent of the first- best revenues, as the problem size is increased. Size is adjusted via a scale factor one to seven. As expected, the complete information bid price scheme performs well, remaining just above 99% despite dropping from an initial value of 99.8%. CBP performs only slightly worse, dipping to a minimum of 97.6% before rising to approximately 1% below the complete information case.

While 1% of potential revenues can represent a substantial value, given the greatly reduced information shared in CBP, its performance level is still surprising. This achievement, as well as that seen in the first two examples, can be attributed to the power of the bid price. An airline’s partner uses all of the information not shared under CBP in determining its own bid prices. As such, much of the additional information shared in the full information case is, to a large degree, redundant.

Figure 6 – Example of Rapid Bid-Price Convergence and Probable Uniqueness under CBP



In figure 6, we look at an example of the convergence of bid prices during the iterative process done at the beginning of each period. Using two extremely different seeds (starting values) – (1) the corresponding full information bid prices, and (2) a zero vector – the bid prices converge to the same values, and they do so in very few iterations. This apparent uniqueness and rapid convergence was seen in many sample alliances that we examined.

5. Prorated CBP (PCPB) Method

In this section, we briefly outline an interesting extension to the CBP model. The Prorated CBP (PCPB) algorithm is to the CBP algorithm what the bid-price proration scheme is to the bid price scheme in the complete information game. It does have some potential benefits, but also some additional requirements and challenges.

5.1. Overview of PCBP Method

As with the basic CBP algorithm described above, the key assumption to decoupling the alliance under PCBP is that an airline's partners' bid prices will remain constant over the remainder of the horizon. This assumption, again, precludes the airline from calculating or forecasting its partner's bid prices, which would be computationally prohibitive.

What differentiates PCBP from CBP is the use of bid-price proration as the agreed upon revenue sharing scheme for the alliance. Unlike the complete information game, here the future share proportions are not calculated directly. Instead, when estimating the revenue shares, the constant bid price assumption provides the airline with a simpler method for estimating the proportions. In period k , we estimate airline c 's share proportion for itinerary j with inventory \bar{x}_c in period $t < k$, given the current posted bid prices for its partners, \bar{p}_k^{-c} , as:

$$\hat{\alpha}_t^{cj}(\bar{x}_c | \bar{p}_k^{-c}) \approx \frac{p_t^{cj}(\bar{x}_c | \bar{p}_k^{-c})}{p_t^{cj}(\bar{x}_c | \bar{p}_k^{-c}) + p_k^{-cj}}$$

where p_t^{cj} is airline c 's "exact" opportunity cost calculated in the single network problem under the PCBP assumption. The adjusted revenue used in airline c 's period k single network problem for an itinerary j request in period $t < k$ is:

$$\bar{R}_t^{cj}(\bar{p}_k^{-c}) = \begin{cases} \hat{\alpha}_t^{cj}(\bar{x}_c | \bar{p}_k^{-c}) R_t^{cj} & j \in N_s \\ R_t^{cj} & j \in N_c \end{cases},$$

The single network problem is then calculated as usual, either exactly or utilizing one of the approximation methods.

Alternately, when using additive bid prices, in period k , we estimate airline c 's share proportion for itinerary j in period $t < k$ as:

$$\hat{\alpha}_t^{cj}(\bar{x}_c | \bar{p}_k^{-c}) \approx \frac{\sum_{i \in A^{cj}} a_{ij} p_t^{ci}(\bar{x}_c | \bar{p}_k^{-c})}{\sum_{i \in A^{cj}} a_{ij} p_t^{ci}(\bar{x}_c | \bar{p}_k^{-c}) + \sum_{i \in A^{-cj}} a_{ij} p_k^{-ci}}$$

where p_t^{ci} is airline c 's "exact" opportunity cost for resource i in period t for the current inventory level and p_k^{-ci} is its partners current (in period k) posted bid price for resource i .

There is, however, an additional step that should be added. Since the bid-price proration scheme is utilized, airline c does actually receive a portion of the surplus revenue from its partner's interline requests. As such, it would be a far less justified approximation to drop that term

for the interline requests to airline c 's partner as we did in reaching equation (4). Instead, the best approach would be to forecast those requests in addition to their own. The benefits and drawbacks of this step are discussed in the next section.

5.2. Advantages and Drawbacks of PCBP

The primary benefit to the PCBP method is the fact that, as noted above, PCBP is based on the bid-price proration scheme, which means that some level of surplus sharing will go on. As discussed by Wright et al (2010), sharing the surplus revenue for a request provides incentive for the partner to hold onto their resources when they may be used for alliance-beneficial (high revenue) requests to the airline in a future period. Whether done directly through actual sharing or indirectly through forecasting, this additional step can be seen as incorporating a higher level of information. CBP does not share this benefit, so its performance suffers in such cases (see section 4 for an example.)

With this benefit comes a drawback: since each the airline receives a portion of the revenue from its partner's interline requests, as noted, it needs to forecast its partner's interline requests. While this requires additional effort to track and estimate those arrival streams, especially since it would only receive the truncated sales information, the computational load of the single network problem would not differ, since the same share proportions are used regardless of which airline serves as the marketing airline. Instead, it would behave as additional demand for the same product.

It should be noted that as unlike CBP, PCBP may generate negative bid prices in the single network problem for certain choices of splitting rules in the case where all bid prices are 0, leading to undefined 0/0 share proportions. With the right choice – e.g., all revenue going to the marketing airline – this problem can be avoided.

6. Summary and Further Research

In this paper, we began exploring the effects of limiting the information shared between alliance partners. Practical, legal, and competitive forces will limit the amount of information that partners are willing and able to share. As such, it is necessary to develop techniques for handling revenue management for interline itineraries under such games of incomplete information.

We provide such a mechanism, utilizing the structure of the bid price sharing scheme and making a constant bid price assumption, a stronger form of the martingale result of Akon and Ata (2009). Under these assumptions, the network is shown to decouple into single-airline network problems, connected by a vector of posted bid prices from each partner. This decoupling allows partners to utilize existing approximation methods for determining the bid prices in their own networks, while handling revenue sharing for interline itineraries.

We discuss the advantages and disadvantages of such a system, and utilize sample networks to illustrate the effects. In general, the scheme is shown to perform nearly as well as and sometimes better than the full-information scheme upon which it is based. These results should be encouraging to practitioners just starting to test such systems.

We finish the paper by describing an extension of our method based on the bid-price proration scheme rather than the bid price scheme. As shown in Wright et al (2010), the bid-price proration scheme can benefit from surplus sharing that the bid price scheme does not. As such, this new scheme has potential benefits in cases where one airline needs incentive to hold onto its inventory for more valuable interline request to its partner. However, some additional complications arise from using this scheme that must be considered when implementing it.

Future research in this area could focus on three particular areas. First, we have considered an extreme level of information sharing, in which airlines know only their partners' bid prices. However, in practice, airlines may actually have more information available to incorporated into their decision processes. Examples include current inventory levels and forecasts of interline request. If, how, and when incorporating this additional information would be beneficial remains an open question.

Second, the CBP scheme decomposes the network into single-network problems. There exists a plethora of mechanisms for determining bid prices for a single airline both in the literature and in use by airlines. It is important to know how the choice of mechanism used by each individual airline affects the performance of CBP or PCBP.

Third, many airlines utilize booking limit (or protection level) inventory control schemes rather than bid price ones. It is possible to create a set of booking limits from a set of bid prices, but it is not clear that the bid prices developed under CBP would be effective in such a control mechanism.

Appendix 1 Proof of Theorems

THEOREM 4.1 (Existence under CBP): In order to demonstrate existence, we must demonstrate that the best response bid prices are continuous in the partner's bid prices and are defined on a compact set. With the assumption of finite upper bounds on the revenue distributions, the bid prices form a compact set, as each is bounded between zero and the maximum potential revenue that can be received for the inventory. Continuity is demonstrated using a proof by induction.

Base Case ($t = 1$):

We know that for all \bar{x}_c and m_{-c} vectors, $\hat{J}_0^c(\bar{x}_c | \bar{p}) = 0$ and, therefore, by extension:

$$\Delta \hat{J}_0^{cj}(\bar{x}_c | \bar{p}) = \hat{J}_0^c(\bar{x}_c | \bar{p}) - \hat{J}_0^c(\bar{x}_c - A^{cj} | \bar{p}) = 0, \quad \forall \bar{x}_c, j, \bar{p}.$$

$$\begin{aligned} \hat{J}_1^c(\bar{x}_c | \bar{p}_k^{-c}) &= \hat{J}_0^c(\bar{x}_c) + \sum_{j \in N} q_1^{cj} E \left[\left(R_1^{cj} - \Delta \hat{J}_0^{cj}(\bar{x}_c | p_k^{-cj}) - p_k^{-cj} \right)^+ \right] \\ &= \sum_{j \in N} q_1^{cj} E \left[\left(R_1^{cj} - p_k^{-cj} \right)^+ \right] \\ &= \sum_{j \in N} q_1^{cj} \left(\int_{p_k^{-cj}}^{\infty} r dF_1^{cj}(r) - \left(1 - F_1^{cj}(p_k^{-cj}) \right) p_k^{-cj} \right) \end{aligned} \quad (10)$$

By assumption, $F_1^{cj}(r)$ is continuous function in r , so $\left(1 - F_1^{cj}(p_k^{-cj}) \right) p_k^{-cj}$ is the product of two continuous functions in \bar{p}_k^{-c} and is, therefore, also continuous in \bar{p}_k^{-c} .

For the first term, we turn to the following proposition that we will also use in the inductive step of the proof:

PROPOSITION A1.1. If $\rho(\bar{p})$ is a continuous nonnegative function in \bar{p} and $F(r)$ is a proper cumulative density function, then:

$$\theta(\bar{p}) = \int_{\rho(\bar{p})}^{\infty} r dF(r)$$

is a continuous function in \bar{p} .

PROOF.

Since $\rho(\bar{p})$ is continuous in \bar{p} , then for each ε , then exists a δ such that $\|\bar{p}_1 - \bar{p}_2\| < \delta$ implies $\|\rho(\bar{p}_1) - \rho(\bar{p}_2)\| < \varepsilon$. Arbitrarily, let $\rho(\bar{p}_1) \geq \rho(\bar{p}_2)$, then because $\rho(\bar{p})$ is nonnegative, we know $\rho(\bar{p}_1) < \varepsilon$. Then,

$$\begin{aligned} \theta(\bar{p}_2) - \theta(\bar{p}_1) &= \int_{\rho(\bar{p}_2)}^{\infty} rdF(r) - \int_{\rho(\bar{p}_1)}^{\infty} rdF(r) \\ &= \int_{\rho(\bar{p}_2)}^{\rho(\bar{p}_1)} rdF(r) \\ &\leq \rho(\bar{p}_1)(F(\rho(\bar{p}_1)) - F(\rho(\bar{p}_2))) \\ &\leq \rho(\bar{p}_1) \\ &< \varepsilon \end{aligned}$$

Since $\theta(\bar{p}_2) - \theta(\bar{p}_1) \geq 0$, for each ε , there exists a δ such that $\|\bar{p}_1 - \bar{p}_2\| < \delta$ implies $\|\theta(\bar{p}_1) - \theta(\bar{p}_2)\| < \varepsilon$. ■

By proposition A1.1, $\int_{\bar{p}_k^{-c}}^{\infty} rdF_1^{cj}(r)$ is continuous in \bar{p}_k^{-c} and, therefore, each term in the summation in (10) is the difference of two continuous functions, so it is continuous. Finally, the whole expression in (10) is the sum of these continuous functions and, thus, $\hat{J}_1^c(\bar{x}_c | \bar{p}_k^{-c})$ and $\Delta \hat{J}_1^{cj}(\bar{x}_c | \bar{p}_k^{-c})$ are continuous in \bar{p}_k^{-c} as well.

Inductive step ($1 < t \leq k$):

Assume (for induction) that $\hat{J}_{t-1}^c(\bar{x}_c | \bar{p}_k^{-c})$ and $\Delta \hat{J}_{t-1}^{cj}(\bar{x}_c | \bar{p}_k^{-c})$ are continuous in \bar{p}_k^{-c}

$\forall \bar{x}_c \geq \bar{0}$, $j \in N_s$. We begin with the decomposition of $\hat{J}_t^c(\bar{x}_c | \bar{p}_k^{-c})$:

$$\begin{aligned}\hat{J}_t^c(\bar{x}_c | \bar{p}_k^{-c}) &= \hat{J}_{t-1}^c(\bar{x}_c | \bar{p}_k^{-c}) + \sum_{j \in N} q_t^{cj} E \left[\left(R_t^{cj} - \Delta \hat{J}_{t-1}^{cj}(\bar{x}_c | \bar{p}_k^{-c}) - p_k^{-cj} \right)^+ \right] \\ &= \hat{J}_{t-1}^c(\bar{x}_c | \bar{p}_k^{-c}) + \sum_{j \in N} q_t^{cj} \left(\int_{\Delta \hat{J}_{t-1}^{cj}(\bar{x}_c | \bar{p}_k^{-c}) + p_k^{-cj}}^{\infty} r dF_1^{cj}(r) - \right. \\ &\quad \left. \left(1 - F_1^{cj}(\Delta \hat{J}_{t-1}^{cj}(\bar{x}_c | \bar{p}_k^{-c}) + p_k^{-cj}) \right) (\Delta \hat{J}_{t-1}^{cj}(\bar{x}_c | \bar{p}_k^{-c}) + p_k^{-cj}) \right)\end{aligned}$$

$\hat{J}_{t-1}^c(\bar{x}_c | \bar{p}_k^{-c})$ is continuous by the inductive assumption. If $F_k^{cj}(r)$ is continuous in r , then second term in the summation is also continuous in p_k^{-cj} . $\Delta \hat{J}_{t-1}^{cj}(\bar{x}_c | \bar{p}_k^{-c}) + p_k^{-cj}$ is the sum of two continuous function and, thus, continuous as well. By Proposition A1.1, the first term in the summation is also continuous and, by similar arguments as the base case, $\hat{J}_t^c(\bar{x}_c | \bar{p}_k^{-c})$ and $\Delta \hat{J}_t^{cj}(\bar{x}_c | \bar{p}_k^{-c})$ are continuous in \bar{p}_k^{-c} . Therefore, $p_k^{cj} = \Delta \hat{J}_{k-1}^{cj}(\bar{x}_c | \bar{p}_k^{-c})$ is continuous in \bar{p}_k^{-c} . ■

THEOREM 4.2 (Uniqueness for special case of CBP): For convergence, we need to show that iterative applications of the transfer prices will converge, which can be shown by:

$$\lim_{n \rightarrow \infty} (Dp_k^c(p_k^{-c}(\bar{p})))^n = 0$$

Where $Df(x)$ is the Jacobian of $f(x)$.

$$\begin{aligned}Dp_k^c(p_k^{-c}(\bar{p})) &= \begin{bmatrix} \nabla p_k^{c1}(p_k^{-c}(\bar{p})) \\ \vdots \\ \nabla p_k^{cn_s}(p_k^{-c}(\bar{p})) \end{bmatrix} \\ &= \begin{bmatrix} \frac{\partial}{\partial p_1}(p_k^{c1}(p_k^{-c}(\bar{p}))) & \cdots & \frac{\partial}{\partial p_{n_s}}(p_k^{c1}(p_k^{-c}(\bar{p}))) \\ \vdots & \ddots & \vdots \\ \frac{\partial}{\partial p_1}(p_k^{cn_s}(p_k^{-c}(\bar{p}))) & \cdots & \frac{\partial}{\partial p_{n_s}}(p_k^{cn_s}(p_k^{-c}(\bar{p}))) \end{bmatrix}\end{aligned}$$

$$\begin{aligned}
&= \begin{bmatrix} \frac{\partial p_k^{c1}(\bar{q})}{\partial q_1} & \dots & \frac{\partial p_k^{c1}(\bar{q})}{\partial q_{n_s}} \\ \vdots & & \vdots \\ \frac{\partial p_k^{cn_s}(\bar{q})}{\partial q_1} & \dots & \frac{\partial p_k^{cn_s}(\bar{q})}{\partial q_{n_s}} \end{bmatrix}^T \cdot \begin{bmatrix} \frac{\partial p_k^{-c1}(\bar{p})}{\partial p_1} & \dots & \frac{\partial p_k^{-c1}(\bar{p})}{\partial p_{n_s}} \\ \vdots & & \vdots \\ \frac{\partial p_k^{-cn_s}(\bar{p})}{\partial p_1} & \dots & \frac{\partial p_k^{-cn_s}(\bar{p})}{\partial p_{n_s}} \end{bmatrix} \\
&= \left(Dp_k^c(\bar{q}) \Big|_{p_k^{-c}(\bar{p})} \right)^T \cdot Dp_k^{-c}(\bar{p}) \\
&= \left(Dp_k^c(p_k^{-c}(\bar{p})) \right)^T \cdot Dp_k^{-c}(\bar{p}) \quad (\text{Alternative form})
\end{aligned}$$

A sufficient condition for this and for uniqueness (contraction mapping) is:

$$\rho(Dp_k^c(p_k^{-c}(\bar{p}))) < 1$$

where ρ is the spectral radius of the matrix. We know that for any induced matrix norm:

$$\begin{aligned}
\rho(Dp_k^c(p_k^{-c}(\bar{p}))) &= \rho\left(\left(Dp_k^c(\bar{q}) \Big|_{p_k^{-c}(\bar{p})} \right)^T \cdot Dp_k^{-c}(\bar{p}) \right) \\
&\leq \left\| \left(Dp_k^c(\bar{q}) \Big|_{p_k^{-c}(\bar{p})} \right)^T \cdot Dp_k^{-c}(\bar{p}) \right\| \\
&\leq \|Dp_k^c(\bar{q})\| \cdot \|Dp_k^{-c}(\bar{p})\|
\end{aligned}$$

Therefore, if we can find a norm for which we can prove $\|Dp_k^c(\bar{p})\| < 1$, then noting that the choice of c , $-c$ is arbitrary, we have proven convergence. We prove this for a special case, using an inductive proof. We begin with the base case for a general network, using the maximum absolute row sum norm, then show the inductive step for the special case (note: the choice of row sum or column sum is arbitrary for the special case since there is only one element.)

Base case ($t = 1$):

We know that for all \bar{x}_c , $\hat{J}_0^c(\bar{x}_c) = 0$ and, therefore, by extension: $\Delta \hat{J}_0^{cj}(\bar{x}_c) = 0, \quad \forall \bar{x}, j$.

$$\begin{aligned}
& \sum_{i \in N_s} \left| \frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_1^{cj}(\bar{x}_c) \right| \\
&= \sum_{i \in N_s} \left| \frac{\partial}{\partial p_k^{-ci}} \left(\hat{J}_1^c(\bar{x}_c) - \hat{J}_1^c(\bar{x}_c - A^{cj}) \right) \right| \\
&= \sum_{i \in N_s} \left| \frac{\partial}{\partial p_k^{-ci}} \left(\begin{aligned} & \hat{J}_0^c(\bar{x}_c) + \left(\sum_{\varphi \in N_c} q_1^{c\varphi} E \left[\left(R_1^{c\varphi} - \Delta \hat{J}_0^{c\varphi}(\bar{x}_c) \right)^+ \right] I_{(\bar{x}_c \geq A^{c\varphi})} \right. \\ & \left. + \sum_{\varphi \in N_s} q_1^{c\varphi} E \left[\left(R_1^{c\varphi} - \Delta \hat{J}_0^{c\varphi}(\bar{x}_c) - p_k^{-c\varphi} \right)^+ \right] I_{(\bar{x}_c \geq A^{c\varphi})} \right) \\ & - \hat{J}_0^c(\bar{x}_c - A^{cj}) - \left(\sum_{\varphi \in N_c} q_1^{c\varphi} E \left[\left(R_1^{c\varphi} - \Delta \hat{J}_0^{c\varphi}(\bar{x}_c - A^{cj}) \right)^+ \right] I_{(\bar{x}_c \geq A^{cj} + A^{c\varphi})} \right. \\ & \left. + \sum_{\varphi \in N_s} q_1^{c\varphi} E \left[\left(R_1^{c\varphi} - \Delta \hat{J}_0^{c\varphi}(\bar{x}_c - A^{cj}) - p_k^{-c\varphi} \right)^+ \right] I_{(\bar{x}_c \geq A^{cj} + A^{c\varphi})} \right) \end{aligned} \right) \right| \\
&= \sum_{i \in N_s} \left| \frac{\partial}{\partial p_k^{-ci}} \left(\begin{aligned} & \sum_{\varphi \in N_c} q_1^{c\varphi} E \left[R_1^{c\varphi} \right] I_{(\bar{x}_c \geq A^{c\varphi})} + \sum_{\varphi \in N_s} q_1^{c\varphi} E \left[\left(R_1^{c\varphi} - p_k^{-c\varphi} \right)^+ \right] I_{(\bar{x}_c \geq A^{c\varphi})} \\ & - \sum_{\varphi \in N_c} q_1^{c\varphi} E \left[R_1^{c\varphi} \right] I_{(\bar{x}_c \geq A^{cj} + A^{c\varphi})} - \sum_{\varphi \in N_s} q_1^{c\varphi} E \left[\left(R_1^{c\varphi} - p_k^{-c\varphi} \right)^+ \right] I_{(\bar{x}_c \geq A^{cj} + A^{c\varphi})} \end{aligned} \right) \right| \\
&= \sum_{i \in N_s} \left| \frac{\partial}{\partial p_k^{-ci}} \left(\sum_{\varphi \in N_c} q_1^{c\varphi} E \left[R_1^{c\varphi} \right] I_{(\bar{x}_c \geq A^{c\varphi})} \tilde{I}_{(\bar{x}_c \geq A^{cj} + A^{c\varphi})} + \sum_{\varphi \in N_s} q_1^{c\varphi} E \left[\left(R_1^{c\varphi} - p_k^{-c\varphi} \right)^+ \right] I_{(\bar{x}_c \geq A^{c\varphi})} \tilde{I}_{(\bar{x}_c \geq A^{cj} + A^{c\varphi})} \right) \right|
\end{aligned}$$

Where $\tilde{I}_{(\bar{x} \geq X)} = 1 - I_{(\bar{x} \geq X)}$ for all non-negative m_c vectors X .

$$\begin{aligned}
&= \sum_{i \in N_s} \left| - \sum_{\varphi \in N_s} q_1^{c\varphi} \left(1 - F_1^{c\varphi} \left(p_k^{-c\varphi} \right) \right) I_{\varphi=i} I_{(\bar{x}_c \geq A^{c\varphi})} \tilde{I}_{(\bar{x}_c \geq A^{cj} + A^{c\varphi})} \right| \\
&= \sum_{i \in N_s} q_1^{ci} \left(1 - F_1^{ci} \left(p_k^{-ci} \right) \right) I_{(\bar{x}_c \geq A^{ci})} \tilde{I}_{(\bar{x}_c \geq A^{cj} + A^{ci})} \\
&\leq \sum_{i \in N_s} q_1^{ci} \\
&\leq 1
\end{aligned}$$

The final inequality holds with a positive probability of *not* having an interline arrival to airline c – this includes a positive probability of no arrival, a local arrival to airline c and/or any arrival to airline $-c$ – as well as with many other conditions.. However, even in the case where equality holds (i.e., airline c is guaranteed to receive a request), the corresponding

norm for airline $-c$ will be 0, because the sum of the probabilities must be 1. So, the condition that the norm of the gradient of the composition of the partners' responses must be less than 1 will be satisfied in period 1 for any game. (Note: the worst case norm of the gradient of the composition of the partners responses is $1/4$ (basically, it is $\leq p(1-p)$ for some p , which is largest when $p = 1/2$.)

Inductive Step ($1 < t < k - 1$):

Inductive assumption:

$$\sum_{i \in N_s} \left| \frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{cj}(\bar{x}_c) \right| < 1 \quad \forall \bar{x}_c \geq \bar{0}$$

Note: we have defined $\hat{J}_k^c(\bar{x}_c) = -\infty$ for $\bar{x}_c < \bar{0}$ and, therefore, $\Delta \hat{J}_k^{cj}(\bar{x}_c) = \infty$ whenever we do not have $\bar{x}_c \geq A^{cj}$. Likewise, for convenience, let:

$$\frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{cj}(\bar{x}_c) = 0 \quad \forall i, j$$

whenever $\bar{x}_c \geq A^{cj}$ does not hold. That is, changing the other prices when there is not sufficient inventory to sell an itinerary does not have an effect on its price.

$$\begin{aligned} & \sum_{i \in N_s} \left| \frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_t^{cj}(\bar{x}_c) \right| \\ &= \sum_{i \in N_s} \left| \frac{\partial}{\partial p_k^{-ci}} \left(\hat{J}_t^c(\bar{x}_c) - \hat{J}_t^c(\bar{x}_c - A^j) \right) \right| \\ &= \sum_{i \in N_s} \left| \frac{\partial}{\partial p_k^{-ci}} \left(\begin{aligned} & \left(\hat{J}_{t-1}^c(\bar{x}_c) + \sum_{\varphi \in N_c} q_t^{c\varphi} E \left[\left(R_t^{c\varphi} - \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c) \right)^+ \right] I_{\bar{x}_c \geq A^{c\varphi}} \right. \\ & \left. + \sum_{\varphi \in N_s} q_t^{c\varphi} E \left[\left(R_t^{c\varphi} - \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c) - p_k^{-c\varphi} \right)^+ \right] I_{\bar{x}_c \geq A^{c\varphi}} \right) \\ & - \left(\hat{J}_{t-1}^c(\bar{x}_c - A^j) - \sum_{\varphi \in N_c} q_t^{c\varphi} E \left[\left(R_t^{c\varphi} - \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c - A^j) \right)^+ \right] I_{\bar{x}_c \geq A^{c\varphi} + A^{cj}} \right. \\ & \left. + \sum_{\varphi \in N_s} q_t^{c\varphi} E \left[\left(R_t^{c\varphi} - \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c - A^j) - p_k^{-c\varphi} \right)^+ \right] I_{\bar{x}_c \geq A^{c\varphi} + A^{cj}} \right) \end{aligned} \right) \right| \end{aligned}$$

$$= \sum_{i \in N_s} \left| \frac{\partial}{\partial p_k^{-ci}} \left(\begin{array}{l} \Delta \hat{J}_{t-1}^{cj}(\bar{x}_c) + \sum_{\varphi \in N_c} q_t^{c\varphi} \left(E \left[\left(R_t^{c\varphi} - \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c) \right)^+ \right] I_{\bar{x}_c \geq A^{c\varphi}} \right. \\ \left. - E \left[\left(R_t^{c\varphi} - \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c - A^{cj}) \right)^+ \right] I_{\bar{x}_c \geq A^{c\varphi} + A^{cj}} \right) \\ + \sum_{\varphi \in N_s} q_t^{c\varphi} \left(E \left[\left(R_t^{c\varphi} - \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c) - p_k^{-c\varphi} \right)^+ \right] I_{\bar{x}_c \geq A^{c\varphi}} \right. \\ \left. - E \left[\left(R_t^{c\varphi} - \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c - A^{cj}) - p_k^{-c\varphi} \right)^+ \right] I_{\bar{x}_c \geq A^{c\varphi} + A^{cj}} \right) \end{array} \right) \right|$$

Using the envelope theorem:

$$= \sum_{i \in N_s} \left(\begin{array}{l} \frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{cj}(\bar{x}_c) + \sum_{\varphi \in N_c} q_t^{c\varphi} \left(\left(1 - F_t^{c\varphi}(\Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c - A^{cj})) \right) \left(\frac{\partial}{\partial p_k^{-ci}} \Delta J_{t-1}^{c\varphi}(\bar{x}_c - A^{cj}) \right) I_{\bar{x}_c \geq A^{c\varphi} + A^{cj}} \right. \\ \left. - \left(1 - F_t^{c\varphi}(\Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c)) \right) \left(\frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c) \right) I_{\bar{x}_c \geq A^{c\varphi}} \right) \\ + \sum_{\varphi \in N_s} q_t^{c\varphi} \left(\left(1 - F_t^{c\varphi}(\Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c - A^{cj}) + p_k^{-c\varphi}) \right) \left(\frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c - A^{cj}) + I_{\varphi=i} \right) I_{\bar{x}_c \geq A^{c\varphi} + A^{cj}} \right. \\ \left. - \left(1 - F_t^{c\varphi}(\Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c) + p_k^{-c\varphi}) \right) \left(\frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c) + I_{\varphi=i} \right) I_{\bar{x}_c \geq A^{c\varphi}} \right) \end{array} \right)$$

For convenience, we let $p_k^{-c\varphi} = 0 \quad \forall \varphi \in N_c$, noting that such itineraries are not interline and, therefore, we do not need to include them in the norm. We also define the following terms, so that we may express the norm in a more usable form.

$$G_t^{c\varphi}(\bar{x}_c, A^{cj}) = \mathbf{max} \left\{ \begin{array}{l} \left(1 - F_t^{c\varphi}(\Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c) + p_k^{-c\varphi}) \right) I_{\bar{x}_c \geq A^{c\varphi}}, \\ \left(1 - F_t^{c\varphi}(\Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c - A^{cj}) + p_k^{-c\varphi}) \right) I_{\bar{x}_c \geq A^{c\varphi} + A^{cj}} \end{array} \right\}$$

$$g_t^{c\varphi}(\bar{x}_c, A^{cj}) = \mathbf{min} \left\{ \begin{array}{l} \left(1 - F_t^{c\varphi}(\Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c) + p_k^{-c\varphi}) \right) I_{\bar{x}_c \geq A^{c\varphi}}, \\ \left(1 - F_t^{c\varphi}(\Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c - A^{cj}) + p_k^{-c\varphi}) \right) I_{\bar{x}_c \geq A^{c\varphi} + A^{cj}} \end{array} \right\}$$

$$\tilde{N}_j = \left\{ \varphi \in N_c \cup N_s : \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c) > \Delta J_{t-1}^{c\varphi}(\bar{x}_c - A^{cj}) \right\}$$

$$\tilde{N}_j = \left\{ \varphi \in N_c \cup N_s : \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c) \leq \Delta J_{t-1}^{c\varphi}(\bar{x}_c - A^{cj}) \right\}$$

\widehat{N}_j is the set of itineraries for which the inventory required becomes less valuable after selling j . We can think of these itineraries as having a complementary relationship. \widetilde{N}_j is the set for which the inventory required becomes more valuable, representing a substitute relationship. $G_t^{c\varphi}(\bar{x}_c, A^{cj})$ will take on the value of the second term in its maximization if and only if $\varphi \in \widehat{N}_j$ and \bar{x}_c is sufficient to fulfill φ plus j . $g_t^{c\varphi}(\bar{x}_c, A^{cj})$ will take on the opposite term.

Resulting in a general expression:

$$= \sum_{i \in N_S} \left(\left(1 - \sum_{\varphi \in N} q_t^{c\varphi} G_t^{c\varphi}(\bar{x}_c, A^{cj}) \right) \frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{cj}(\bar{x}_c) \right. \\ \left. + \sum_{\varphi \in N} q_t^{c\varphi} \begin{pmatrix} G_t^{c\varphi}(\bar{x}_c, A^{cj}) \\ -g_t^{c\varphi}(\bar{x}_c, A^{cj}) \end{pmatrix} \begin{pmatrix} \left(I_{\varphi \in \widehat{N}_j} + I_{\varphi \in \widetilde{N}_j} \tilde{I}_{\bar{x}_c \geq A^{cj} + A^{c\varphi}} \right) \begin{pmatrix} \frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{cj}(\bar{x}_c) \\ -\frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c) - I_{\varphi=i} \end{pmatrix} \\ + \left(I_{\varphi \in \widehat{N}_j} I_{\bar{x}_c \geq A^{cj} + A^{c\varphi}} \right) \begin{pmatrix} \frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{cj}(\bar{x}_c) \\ + \frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c - A^{cj}) + I_{\varphi=i} \end{pmatrix} \end{pmatrix} \right) \\ \left. + \sum_{\varphi \in N} q_t^{c\varphi} g_t^{c\varphi}(\bar{x}_c, A^{cj}) \frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{cj}(\bar{x}_c - A^{c\varphi}) \right)$$

The final term comes from noting the following relationship:

$$\frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{cj}(\bar{x}_c) + \frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c - A^{cj}) - \frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c) = \frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{cj}(\bar{x}_c - A^{c\varphi})$$

Using the triangle inequality of absolute values:

$$\begin{aligned}
& \left(\left(1 - \sum_{\varphi \in N} q_t^{c\varphi} G_t^{c\varphi}(\bar{x}_c, A^{c\varphi}) \right) \sum_{i \in N_S} \left| \frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{cj}(\bar{x}_c) \right| \right. \\
& \leq \left. + \sum_{\varphi \in N} q_t^{c\varphi} \left(\begin{array}{l} G_t^{c\varphi}(\bar{x}_c, A^{c\varphi}) \\ -g_t^{c\varphi}(\bar{x}_c, A^{c\varphi}) \end{array} \right) \left(\begin{array}{l} \left(I_{\varphi \in \tilde{N}_j} + I_{\varphi \in \tilde{N}_j} \tilde{I}_{\bar{x}_c \geq A^{c\varphi} + A^{c\varphi}} \right) \sum_{i \in N_S} \left| \frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{cj}(\bar{x}_c) \right| \\ - \left(\frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c) + I_{\varphi=i} \right) \end{array} \right) \right. \\
& \left. + \left(I_{\varphi \in \tilde{N}_j} I_{\bar{x}_c \geq A^{c\varphi} + A^{c\varphi}} \right) \sum_{i \in N_S} \left| \frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{cj}(\bar{x}_c) \right| \right. \\
& \left. + \left(\frac{\partial}{\partial p_k^{-ci}} \Delta \hat{J}_{t-1}^{c\varphi}(\bar{x}_c - A^{c\varphi}) + I_{\varphi=i} \right) \right) \\
& + \sum_{\varphi \in N} q_t^{c\varphi} g_t^{c\varphi}(\bar{x}_c, A^{c\varphi}) \sum_{i \in N_S} \frac{\partial}{\partial p_k^{-ci}} \left| \Delta \hat{J}_{t-1}^{cj}(\bar{x}_c - A^{c\varphi}) \right| \right) \quad (11)
\end{aligned}$$

The weights of the three terms will sum to 1 and the first and last terms are <1 by the inductive assumption. Unfortunately, as demonstrated in §4.4.1, for the general case, the middle term cannot be shown to be <1 . We can, however, show this for the special case.

In this special case, each airline operates one leg and sells two itineraries: one intraline and the interline itinerary combining the two legs. Without loss of generality, let φ be the intraline itinerary for airline c and j be the interline itinerary. Notice that in this case, the terms in (11) are for $i = j$ only and the inductive assumption reduces to:

$$\left| \frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{cj}(x) \right| < 1 \quad \forall x \geq 0$$

There is a single leg, so the inventory for c can be expressed as a single value, x . From Talluri and van Ryzin (1998), we know that $j \in \tilde{N}_j$. However, since they utilize the exact same inventory, we know that not only is $\varphi \in \tilde{N}_j$, but:

$$\frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{c\varphi}(x) = \frac{\partial}{\partial p_k^{-cj}} \left(\hat{J}_{t-1}^c(x) - \hat{J}_{t-1}^c(x-1) \right) = \frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{cj}(x)$$

Therefore, the expression in (11) reduces to:

$$\begin{aligned}
& \left(\begin{aligned}
& \left(1 - q_t^{c\varphi} \left(1 - F_t^{c\varphi} \left(\Delta \hat{J}_{t-1}^{c\varphi}(x) \right) \right) \right) \left| \frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{cj}(x) \right| \\
& - q_t^{cj} \left(1 - F_t^{cj} \left(p_k^{-cj} + \Delta \hat{J}_{t-1}^{cj}(x) \right) \right) \left| \frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{cj}(x) \right| \\
& + \left(q_t^{c\varphi} \left(1 - F_t^{c\varphi} \left(\Delta \hat{J}_{t-1}^{c\varphi}(x) \right) \right) \right) \left| \frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{cj}(x) \right| - \left(\frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{c\varphi}(x) \right) \left| \frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{cj}(x) \right| \\
& - q_t^{c\varphi} \left(1 - F_t^{c\varphi} \left(\Delta \hat{J}_{t-1}^{c\varphi}(x-1) \right) \right) \left| \frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{cj}(x) \right| \\
& + \left(q_t^{cj} \left(1 - F_t^{cj} \left(p_k^{-cj} + \Delta \hat{J}_{t-1}^{cj}(x) \right) \right) \right) \left| \frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{cj}(x) \right| - \left(\frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{cj}(x) + 1 \right) \left| \frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{cj}(x) \right| \\
& - q_t^{cj} \left(1 - F_t^{cj} \left(p_k^{-cj} + \Delta \hat{J}_{t-1}^{cj}(x-1) \right) \right) \left| \frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{cj}(x) \right| \\
& + \left(q_t^{c\varphi} \left(1 - F_t^{c\varphi} \left(\Delta \hat{J}_{t-1}^{c\varphi}(x-1) \right) \right) \right) \left| \frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{cj}(x-1) \right| \\
& + q_t^{cj} \left(1 - F_t^{cj} \left(p_k^{-cj} + \Delta \hat{J}_{t-1}^{cj}(x-1) \right) \right) \left| \frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{cj}(x-1) \right|
\end{aligned} \right) \\
& = \left(\begin{aligned}
& \left(1 - q_t^{c\varphi} \left(1 - F_t^{c\varphi} \left(\Delta \hat{J}_{t-1}^{c\varphi}(x) \right) \right) \right) \left| \frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{cj}(x) \right| \\
& - q_t^{cj} \left(1 - F_t^{cj} \left(p_k^{-cj} + \Delta \hat{J}_{t-1}^{cj}(x) \right) \right) \left| \frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{cj}(x) \right| \\
& + \left(q_t^{cj} \left(1 - F_t^{cj} \left(p_k^{-cj} + \Delta \hat{J}_{t-1}^{cj}(1) \right) \right) \right) \left| -1 \right| \\
& - q_t^{cj} \left(1 - F_t^{cj} \left(p_k^{-cj} + \Delta \hat{J}_{t-1}^{cj}(x-1) \right) \right) \left| -1 \right| \\
& + \left(q_t^{c\varphi} \left(1 - F_t^{c\varphi} \left(\Delta \hat{J}_{t-1}^{c\varphi}(x-1) \right) \right) \right) \left| \frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{cj}(x-1) \right| \\
& + q_t^{cj} \left(1 - F_t^{cj} \left(p_k^{-cj} + \Delta \hat{J}_{t-1}^{cj}(x-1) \right) \right) \left| \frac{\partial}{\partial p_k^{-cj}} \Delta \hat{J}_{t-1}^{cj}(x-1) \right|
\end{aligned} \right) \\
& < 1
\end{aligned}$$

The final inequality holds because the weights sum to 1 and the absolute values in the first and third terms are < 1 due to the inductive assumption. \blacksquare

Acknowledgements

The author would like to thank Professors Robert Shumsky and Harry Groenevelt for their guidance and support. I appreciate comments from the numerous researchers, analysts and revenue managers who have provided me with feedback, as well as information about the current state of revenue management in alliances.

References

- Akan, M., and B. Ata. 2009. Bid-Price Controls for Network Revenue Management: Martingale Characterization of Optimal Bid Prices. *Mathematics of Operations Research*. **34** (11) 912-936.
- Boyd, E. A. 1998a. Airline Alliances. *OR/MS Today*. October 28-31.
- Boyd, A. E. 1998b. Airline Alliance Revenue Management. PROS Strategic Solutions Technical Report, PROS Revenue Management, 3100 Main Street, Suite #900, Houston, TX 77002.
- Bureau of Transportation Statistics, U.S. Government (BTS). 2006. Airline Origin and Destination Survey (DB1B). Available at <http://www.transtats.bts.gov/DataIndex.asp>.
- Fernandez de la Torre, P. E. 1999. Airline alliances: The airline perspective. Masters thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Fudenberg, D. and Tirole, J. 1991. *Game Theory*. MIT Press.
- Isler, K. and Imhof, H. 2008. A Game Theoretic Model for Airline Revenue Management and Competitive Pricing. *Journal of Revenue and Pricing Management*. **7** 384-396.
- Kunnumkal, S. and Topaloglu, H. 2010. Computing Time-Dependent Bid Prices in Network Revenue Management. *Transportation Science*. **44** (1) 1526-1547.
- McGill, J. and van Ryzin, G. 1999. Revenue Management: Research Overview and Prospects. *Transportation Science*. **33** (2) 233-256.
- Netessine, S. and Shumsky, R. 2005. Revenue Management Games: Horizontal and Vertical Competition. *Management Science*. **51** (5) 813-831.
- Raza, A. and Akgunduz, A. 2008. An Airline Revenue Management Pricing Game with Seat Allocation. *International Journal of Revenue Management*. **2** (1) 42-62.
- Song, J. 2007. Game theoretic revenue management models for hotel room inventory control. *ETD Collection for McMaster University*. Paper AAINR36048.

- Shumsky, R. A. 2006. The Southwest Effect, Airline Alliances, and Revenue Management. Forthcoming in the *Journal of Revenue and Pricing Management*.
- Talluri, K. and van Ryzin, G. 1998. An Analysis of Bid-Price Controls for Network Revenue Management. *Management Science*. **44** (11) 1577-1593.
- Talluri, K. and van Ryzin, G. 2004 *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers.
- Topaloglu, H. 2009. Using Lagrangian Relaxation to Compute Capacity-Dependent Bid Prices in Network Revenue Management. *Operations Research*. **57** (3) 637-649.
- Vinod, B. 2005. Alliance Revenue Management. *Journal of Revenue and Pricing Management* (Practice Papers), **4**(1) 66-82.
- Vulcano, G., van Ryzin, G. and Maglaras, C. 2002. Optimal Dynamic Auctions for Revenue Management. *Management Science*. **48** (11) 1388-1407.
- Wynne, R. 1995. From Queen Airs to RJs: Evolution in Major-Regional Airline Alliances and Code Sharing. *Handbook of Airline Economics*. Aviation Week Group, McGraw-Hill, New York.
- Wright, C., Shumsky, R. and Groenevelt, H. 2010. Dynamic Revenue Management in Airline Alliances. *Transportation Science*, **44**(1) 15-37.