Exchanging values of unsold seats with airline alliance partners in a competitive environment

Alyona Michel[1,2] and Peter P. Belobaba[1]
[1]Massachusetts Institute of Technology
Cambridge, Massachusetts
[2]**one**world Airline Alliance
New York, New York
Emails: almichel@mit.edu and belobaba@mit.edu

***Abstract*—Airlines in alliances expand the scope of their networks through code sharing, which allows an airline to sell seats on another airline's flight as if the seats were its own. However, code sharing is problematic for airline revenue management (RM) systems. RM improves revenues by placing booking limits on the availability of lower fares in a market. Alliance partners do not work together to determine seat availability for code share flights, producing asymmetric booking limits. Thus, seat inventory is sold for less than its value to the network, and revenues are lower. In this paper, we test whether alliance partners facing competition from another alliance benefit from using information about the estimated value of unsold seats on partner flights in their RM systems. It is found that alliance network characteristics and the behavior of the competition affect the performance of the two methods tested. Results from the simulation show that under favorable circumstances, benefits reach 0.43% if the alliance carries a large share of code share traffic at high fares and does not risk losing non-code share high fare traffic. Depending on network characteristics, the sophistication of the competition's RM can decrease and even reverse revenue benefits.**

I. INTRODUCTION

In the last 2 years, airline alliances have acquired a long list of new members, and few of the world's remaining quality airlines—large or small—have not either joined an alliance or at least partner with other airlines. In the year 2013, alliances earned 67% of worldwide passenger airline revenue, and earned a revenue premium on the traffic they carried [14]. With alliances constantly searching for new members (or code share partners) to fill the gaps in their worldwide networks, they will continue to grow in size and importance.

The three major alliances, **one**world, Star and SkyTeam, complete intensely on the world's long-haul business routes. A major feature of alliance cooperation in providing good service in an important market is code sharing. Code sharing allows an alliance partner to sell seats on another airline's flight as if they were flown by the selling partner. This allows partners to expand the scope of their networks to distant destinations they do not (or cannot) fly to. Thus, partners "virtually" connect distant points, serving global markets each would otherwise be unable to serve without having flying rights and deploying its own planes and personnel.

Airlines use revenue management (RM) to attempt to optimize the revenue earned from available seats on the flights in their network. RM systems use availability control to place booking limits on the availability of certain fares in a market, or origin-destination pair (O-D). Restricting the sale of low fares later in the booking process (time window during which tickets are sold) improves revenues because tickets are sold at higher fares when customers have a higher willingness to pay. In the absence of alliances, an airline's traffic consists of local traffic (when the passenger flies on one flight leg from his origin to his destination) and connecting (when he flies on multiple flight legs). An airline's RM system only has access to information about the state of the airline's own network. Thus, airlines cannot accurately evaluate booking limits for code share itineraries (which consist of flight legs operated by different partners) due to the lack of information about partners' operated flight legs.

Whether they are unable or unwilling, alliance partners do not work together to determine booking limits on code share flights. This often produces asymmetric booking limits on itineraries that involve code share flights. Thus, seat inventory is sold for less than its potential value to the network, and revenues are lower for the alliance. In this paper, we test whether it is beneficial for revenues if alliance partners facing competition from another alliance incorporate information into their RM system's decision-making process about the estimated value of unsold seats on their partner's operated flights, called "bid prices".

Two methods are examined. Bid price sharing (BPS) incorporates the bid prices into the decision process after booking limits have been determined for the airline's own local and connecting itineraries. Dynamic valuation (DV) incorporates bid prices into the calculation of booking limits at the same time as other booking limits are determined, having a larger effect on booking limits on the airline's own local and connecting fares.

The goal of this paper is to test which method performs better for alliances with a network structure resembling that of the three global alliances, which have strong partnerships across the Atlantic. This paper addresses a potential research direction that follows from a previous paper [17], which tested the effects of BPS in a US-based alliance network with a very different structure. Also, this paper aims to gain insights into the performance of the two methods depending on the behavior of the competing alliance, and the network characteristics of the alliances.

II. BACKGROUND

This section provides background on alliance RM concepts important to the experiments in this paper.

A. Three Types of Alliance Flight Traffic: Local, Connecting, and Code Share

Figure 1 below illustrates three flight legs. Alliance partner 1 operates the flights DEN-ORD and ORD-FRA, and also code shares on partner 2's operated flight FRA-BUD. Partner 2 operates FRA-BUD and also code shares on partner 1's flight ORD-FRA. Various combinations of these flights comprise the three components of alliance flight traffic:

In an ***own local*** itinerary, the passenger flies on one flight leg from his origin to his intended destination. The figure illustrates three possible local itineraries, DEN-ORD, ORD-FRA, and FRA-BUD.

In an ***own connecting*** itinerary, the passenger flies on multiple flight legs of a single airline from his origin to his intended destination. In this illustration, only one own connecting itinerary is possible: DEN-ORD-FRA on partner 1.

In a ***code share itinerary***, a passenger flies on multiple flight legs operated by different partners, but marketed under a single airline's code, i.e., sold by that airline. In this illustration, a passenger may purchase the code share itinerary ORD-FRA-BUD from either partner 1 or partner 2, because both partners sell the itinerary under their operating codes.

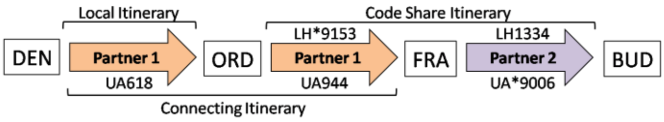

Figure 1. Components of Alliance Flight Traffic.

B. Alliance Revenue Management

Alliance RM consists of the following four components, described below: recording and forecasting, code share valuation, seat allocation (or optimization), and code share availability control. This paper's assumptions regarding the RM characteristics in the experiments are also stated.

An airline's RM system draws on a historical database of prior bookings to forecast future demand by fare class and O-D market. ***Recording and forecasting*** of bookings refers to how airlines record accepted bookings for later use in forecasting. Full information of the code share booking is needed from the booking agent (such as a Global Distribution System), in order to know the full itinerary and distinguish it from a local booking when making future forecasts. This paper assumes that alliance partners record the full itinerary and forecast code shares separately from locals.

An airline sets fares (which range from low to high, depending on the booking class) for the local and connecting itineraries it offers for sale. These fares are provided as input to the RM system (along with demand forecasts) for calculating booking limits. Less straightforward, however, is the value an airline should attribute to a code share itinerary. ***Code share valuation*** refers to the value used as input to the RM seat allocator optimization model. The possibilities include using the total fare of the code share itinerary, the local fare of the partner's own traversed flight leg (both total and local valuation overvalue the code share itineraries), a proration of the total fare, or the novel method of DV. In this experiment, we use local valuation of code share itineraries, as representative of current airline practices, for all tests that do not involve dynamic valuation. Figure 2 presents an example of the valuation of ORD-FRA-BUD. In this example, local valuation values the itinerary at $700, whereas DV values it at $650, both higher than the O-D fare.

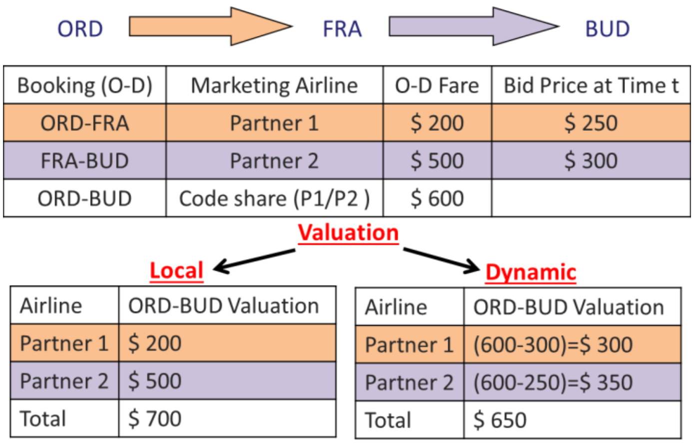

Figure 2. Example of code share local/dynamic valuation.

Airlines then determine how many seats on a flight leg to make available for sale. Seat availability is restricted through booking limits on low fares in an O-D market. Booking limits are calculated during ***seat allocation*** (or optimization). The RM system may use optimization that is leg-based (optimizing revenues on an individual flight leg) or O-D/network based (aiming to optimize revenues over the entire network by considering flight leg opportunity costs). ***Expected Marginal Seat Revenue (EMSRb)*** developed by Belobaba [1, 2, 3], is the leg-based method used in this paper. The heuristic calculates the expected revenue from an empty seat on a flight leg and allows a booking to be made if the fare exceeds the EMSR.

The two network-based methods used in this paper are ***Displacement Adjusted Virtual Nesting (DAVN)*** and ***Probabilistic Bid Price (ProBP)***. In DAVN, a network linear program produces shadow prices that represent the opportunity costs of filling the last empty seat on a flight leg (for more information, see [22]). Note that, if the airline flight leg is not booked to near capacity, most shadow prices produced by the deterministic linear program of DAVN are zero. However, as bookings on the flight leg reach capacity, these shadow prices quickly approach an expensive fare on that flight leg. The RM optimizer uses the total O-D fare less a flight leg's shadow price to determine which O-D fares involving a flight leg to make available for sale. If a leg shadow price is small, the estimated network value of selling the seat will be higher. Thus, a seat on that flight leg is more likely to be available for booking at a low fare. Because it uses multiple fare level ranges (called "buckets") to determine open/closed booking availability, DAVN is a robust network RM method, with results less sensitive to small variability in historical input data or demand forecasting methods. However, it is still a heuristic.

ProBP optimally prorates the total O-D fare to the affected flight legs using an iterative algorithm. The values of the last and marginal seats on a flight leg are calculated. These EMSR values are prorated to the flight legs comprising the itinerary until convergence is reached (more information in [8]). The resulting prorated EMSR values are the "bid prices" of the affected flight legs. Theoretically more optimal than DAVN, this method is also more sensitive to inputs and forecasts. It can perform worse than DAVN, especially at low demand levels or when there is high demand variability. This paper tests using BPS or DV when the alliance partners both use symmetric DAVN or ProBP, and when the competing alliance uses EMSR, identical RM, or identical RM with BPS or DV.

RM methods all produce some estimate of the value of an unsold seat, which we refer to, without loss of generality, as "bid prices". Figure 3 below summarizes the characteristics of the three RM methods tested in this paper, and the types of bid prices generated.

RM System	Level of Forecasting	Level of Optimization	Level of Control	Control Mechanism	Bid Price Type
Leg EMSRb	Leg	Leg Optimization	Leg	Booking Limits	EMSRc value
DAVN	Path	Network Optimization, Leg Capacity Constraints	Leg	Booking Limits	LP Shadow Price
ProBP	Path	Network Optimization	Path	Bid Prices (Converged prorated fares)	

Figure 3. Characteristics of the RM methods EMSRb, DAVN, and ProBP.

The RM system must decide which code share itineraries are available at which fare, referred to as ***code share availability control***. These include bid price sharing control (proposed by [9]) and standard availability control, when code share booking limits are determined inside the RM optimizer, during the seat allocation step. Importantly, the code share valuation inputs may differ. If local valuation is used, code shares are "worth as much" to the network as own local bookings. With standard control using DV, code shares are treated as distinct from locals because they receive a different valuation, which may be lower or higher than the local valuation. With BPS, standard availability control with local valuation is used to calculate bid prices. Bid prices are then shared by partners to inform seat availability decisions after booking limits have been determined for own local and connecting itineraries [15].

C. Description of the Passenger Origin-Destination Simulator

The Passenger Origin-Destination Simulator (PODS) is a software simulation developed at The Boeing Company in the mid-1990s by Hopperstad, Berge and Filipowski [13]. It is used in this research to test the effects of sharing seat value information among alliance partners. Two underlying models within PODS represent the processes of passenger choice and airline RM systems, which interact at the levels of seat allocation and passenger choice/decision [6]. The software simulates the booking process into 16 time frames covering 63 days prior to flight departure. At the start of each time frame, seat allocation is reoptimized based on bookings to date. Each simulation trial consists of 600 weeks, or samples. To eliminate bias from any initial conditions, the first 200 samples are discarded. An average of the remaining 400 samples gives the simulation trial output, and ensures small standard deviation of the results. The PODS Research Consortium at MIT uses the software to test the effects of different revenue management techniques on airlines. Participation from 9 airlines gives the Consortium a perspective that is oriented to the realities and operating constraints faced by airlines today.

D. Prior Work on Alliance Revenue Management

Booking limits on code share itineraries are sub-optimal because the seats in the network are not being allocated as if they were part of one network, and each partner does not have full knowledge about the state of the other's system. One simulation result estimates that network RM should improve revenues by 1.0-1.8%, but if code share traffic is present, the gains drop to 0.5-0.7% [14]. This loss can be exacerbated because code shares are often long-haul international routes with high revenue contribution. Therefore, prior research in the field of operations research has attempted to address the problem of alliance revenue management and propose optimal solutions [7, 12, 15, 17, 18, 20, 21, 23, 24].

Some solution approaches involve modifying the valuation of code share itineraries [15, 17, 18, 23]. Findings in [23] show that the performance of static valuation schemes is nearly as good as the best dynamic schemes (achieving about 90% of revenues), but that static schemes do not adjust well as network parameters change, which could rapidly reduce revenues. The network in [23] consists of just two flight legs with ten seats each, in which the proportion of code share itineraries was varied. In [15, 17, 18], it is found that DV using dated partner bid prices performs slightly better than BPS in a complex US-based network with constant network parameters.

An option that could achieve optimal joint revenues, suggested by both [7] and [21], is to exchange seats among the alliance partners until the relative values of the seats to each airline's network are equal. After the seats are exchanged and paid for among the carriers, each airline has individual control over the seats by its own RM system. In [7], this proposition is formulated in terms of marginal seat values, and in [21] in terms of bid prices, but the idea is the same. In the proposed scenario, the resource allocation is optimal and expected revenue for the alliance can be maximized. However, in practice, airlines do not calculate bid prices or expected seat revenues for seats on flights that they do not operate, as they do not have access to the booking histories and forecasts for those non-operated flights. Significant technological changes to the airlines' RM systems and antitrust immunity may be required before such a scheme could be implemented.

Applying finance theory, [12] proposes an options-based approach to capacity control on a single flight leg, where the marketing carrier can purchase options from the operating carrier for the right to buy seat inventory in the future at a pre-determined strike price. No theoretical basis or evidence is presented that this scheme performs better than the other methods proposed in the literature. In [20], a deterministic LP model inspired by that in [22] is developed, similar to the approach of DAVN. The LP is solved over the joint alliance network, and then the constraints linking the partners' decisions to the joint alliance network solution are relaxed through the dual prices, thus decomposing the problem into smaller problems by airline.

Under BPS control, a code share itinerary is available for booking if the fare exceeds the sum of the bid prices on all the legs traversed. It was proposed in [9] that airlines either exchange either actual bid prices (if they have the legal ability to do so), or to infer bid prices from the lowest available fares, an idea re-iterated in [21] in order to inform one another of the value of seat inventory on code share flight legs. Bid price inference may be a feasible option for airlines that do not have the antitrust immunity to share such information. The small-scale simulations in [23] show that bid price control produces gains just as good as those of the dynamic transfer price schemes in the cooperative game of complete information.

III. Methods and Experiment Setup

A. Characteristics of PODS Alliance Network E

Network E represents two competing alliances, each with partner hubs across the Atlantic. A schematic of network E is illustrated in Figure 4. The two competing alliances have one set of partners (airlines 1 and 2) whose hubs are located in the central United States, while the other two partners (airlines 3 and 4) have hubs located in Europe. Airlines 1 and 3 form alliance 1, and airlines 2 and 4 make up alliance 2. The spoke cities emanate from the continental hubs, with 10 in the northern part and 10 in the southern part of Europe, as well as 10 in the western part and 10 in the eastern part of the United States. In the baseline case for network E, all airlines use EMSRb with leg forecasting.

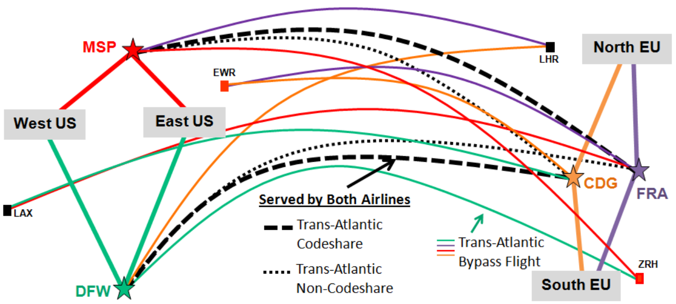

Figure 4. Trans-Atlantic Network E.

Roughly speaking, network E has a dumbbell structure. The hub-to-hub trunk routes (served by large aircraft) carry a large amount of passengers across the Atlantic into the hubs and feed the smaller hub-to-spoke routes. Such trunk flights occur three times a day, providing connecting opportunities during three banks. The routes are served by each airline, resulting in a total of six hub-to-hub flights and 12 additional trans-Atlantic flights per bank. In addition, some local and connecting traffic does not cross the Atlantic, but traverses spoke-to-hub and hub-to-spoke routes, staying on the same continent. Also, there are hub-bypass routes crossing the Atlantic that go from a continental hub to a major city on the other continent, without requiring a connection at the partner hub. The different types of routes ensure that the network is varied and representative of current alliance networks.

B. Bid Price Sharing for Code share Availability Control

BPS refers to any process of partners exchanging network displacement cost information, or the estimated value of selling one item of seat inventory on a flight leg. If the seat allocation method used by the RM system is DAVN, then the resulting bid prices are actually shadow prices from a deterministic linear program for each leg on an airline's own network. The total code share fare (CSfare) must exceed the sum of the minimum threshold of the lowest open fare range "bucket" and the partner's shadow price (SP). This must be satisfied for both partners, known as "dual control". The BPS availability control equation is:

$$CSfare \geq Bucket_{\min}{}^{t} + SP_{partner}{}^{t}. \tag{1}$$

With PROBP, iterative optimal proration of fares produces bid prices for each leg. The availability control equation upon BPS is:

$$CSfare \geq BP_{own}{}^{t} + BP_{partner}{}^{t}. \tag{2}$$

Figure 5 shows an illustration of the BPS controls using DAVN or ProBP.

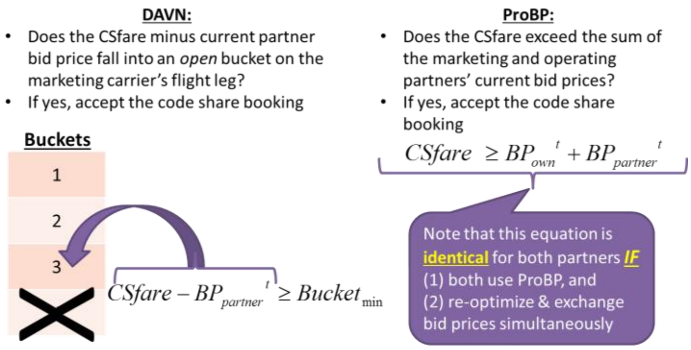

Figure 5. DAVN and ProBP Code Share Control Comparison.

In the above equations, the partner's bid price is the most recent available. The experiments described in this paper use daily optimization for ProBP and daily calculation of shadow prices for DAVN (though booking limits on fare buckets are recalculated at the start of each of 16 time frames). Prior experiments with reoptimization every time frame and every 200 bookings are described in [17] and [18]. As part of the experiments described in this paper, less frequent optimization was also tested. Exactly as before, it was found that more frequent optimization slightly improves results for ProBP, and largely for DAVN. The bid prices more accurately reflect the state of the network, thus the RM systems produce more optimal booking limits and more revenue is generated. Because the results were so similar, this paper only shows the results using daily reoptimization.

This paper thus assumes that airlines have the technological capability to reoptimize, exchange bid prices, and incorporate partner bid prices into their RM system (with DV) on a daily basis. It is acknowledged that a range of capabilities exist among today's airlines, some with integrated systems and real time communication among parent airlines and their subsidiaries, and more rudimentary communication between other alliance partners. BPS and DV require sharing information so quickly and frequently that they may require more advanced RM systems and communication technology on the part of the alliance members.

C. BPS Availability Control vs Dynamic Valuation

The key difference between BPS and DV is that BPS incorporates the bid prices into the decision process after booking limits have been determined for the partner's own local and connecting itineraries. DV incorporates bid prices into the calculation of booking limits at the same time as other booking limits are determined, having a larger effect on booking limits on own local and connecting fares.

With DV, the valuation of code shares is dynamic because it is changing with every RM reoptimization. Recall that with BPS, the value of a code share itinerary used as input to partner 1's RM optimizer is the same as a local online fare of a passenger flying only partner 1's flight legs. In DV, the code share valuation used as input to the optimizer (CSvaluation) is the total code share fare (CSfare) minus the most recent bid price from partner 2

$$CSvalution = CSfare_{own}^{\;t} - BP_{partner}^{\;t}. \qquad (3)$$

This resulting valuation can be lower or higher than local valuation, depending on the bid prices and code share fare. Refer to figure 2 for an example of dynamic valuation. The results using DV in [15] are very encouraging and show revenue gains near to those of BPS for code share availability control in DAVN systems, and much better results in ProBP (about a 0.25% revenue gain for DAVN and ProBP). It is argued that improved performance of ProBP under dynamic valuation is a result of the sensitivity of the ProBP optimization to fare inputs. This paper continues the examination of dynamic valuation in a different network setting, with the exchange of current bid prices and more frequent optimization from that used in [15], with results indicating more complicated conclusions.

D. Dimensions Tested in Experiment

The focus of this paper is alliance revenue management in a competitive environment, which previous research has not addressed. We also test the benefits of BPS or DV when the alliance uses the differing network RM methods of DAVN and ProBP. Although tests were performed to compare daily optimization with optimization once per time frame, the results are similar to prior findings that show that daily optimization largely benefits DAVN, and are not shown. The alliance may face a competing alliance that uses either EMSR, the same network RM method, or the same network RM method in combination with BPS or DV. Two different demand levels are tested to understand how the demand level affects the performance of each method. The aim is to gain insights into the performance of BPS compared with DV in different competitive and network environments. Table 1 provides a summary of the dimensions tested in the experiment.

TABLE I. SUMMARY OF EXPERIMENT DIMENSIONS

Experiment:	**1**	**2**	**3**	**4**	**5**	**6**
Alliance 1	DAVN	ProBP	EMSRb		DAVN	ProBP
Alliance 2	EMSRb		DAVN	ProBP	DAVN	ProBP
Code share valuation	Local and Dynamic Valuation					
Availability Control	Standard control with local or DV, or Bid Price Sharing					
BPS Frequency	Daily exchange of bid prices					
Demand Level	High (86-88%) and medium (82-83%) load factor					

IV. RESULTS

A. Alliance 1 Uses BPS or DV, Alliance 2 Does Not

Figure 6 shows the results obtained when alliance 1 uses BPS or DV under different demand levels. Under medium demand, DV leads to the highest revenue gains for both DAVN and ProBP. However, ProBP with BPS performs better than DV at high demand. When the competitor also uses ProBP, the pattern of gains over standard ProBP is similar, but smaller.

With DAVN for alliance 1 and EMSRb for the competitor, BPS does not improve revenues at the medium demand level. At high demand, gains occur from both BPS and DV. When the competitor uses DAVN as well, the order of gains from BPS and DV remains the same, but all methods produce revenue gains at both medium and high demand levels because of a beneficial competitive feedback effect.

Alliance1	Alliance2	Demand	RM only	ΔBPS	ΔDV
ProBP	EMSRb	Medium	1.48	0.25	0.32
		High	4.66	0.43	0.33
ProBP	ProBP	Medium	0.72	0.12	0.18
		High	1.94	0.31	0.23
DAVN	EMSRb	Medium	1.62	-0.02	0.06
		High	4.53	0.23	0.32
DAVN	DAVN	Medium	1.00	0.12	0.11
		High	2.23	0.11	0.15

Figure 6. Alliance 1 uses BPS or DV, Alliance 2 does not.

The revenue gains are still present for alliance 1 even when alliance 2 uses network RM as its revenue management method rather than EMSRb. However, the resulting revenue gains are smaller when the competitor uses a more intelligent RM system.

B. Alliance 2 Uses BPS or DV, Alliance 1 Does Not

Figure 7 below shows that alliance 2, using ProBP at any demand level, and DAVN at high demand, also benefits from BPS when the competitor uses EMSR, but less than the benefit for alliance 1. When the competition uses network RM, the revenue gains are generally smaller than they were for alliance 1 as well. Also, note that alliance 2 benefits less from network RM in general, compared with alliance 1, due to the fact that it carries a higher proportion of local traffic than alliance 2. Alliance 2 benefits less from DV than BPS, the opposite of the effect on alliance 1.

Alliance1	Alliance2	Demand	RM only	ΔBPS	ΔDV
EMSRb	ProBP	Medium	1.59	0.07	0.08
		High	4.24	0.37	0.08
ProBP	ProBP	Medium	0.64	0.08	0.00
		High	1.63	0.38	0.14
EMSRb	DAVN	Medium	1.51	-0.09	-0.03
		High	4.20	0.19	0.02
DAVN	DAVN	Medium	0.58	-0.06	-0.06
		High	1.54	0.16	-0.02

Figure 7. Alliance 2 uses BPS or DV, Alliance 1 does not.

C. Both Alliances Use BPS or DV, High Demand Only

Figure 8 shows the results obtained when both alliances use BPS or DV, at high demand only. Alliance 1 retains a revenue gain from BPS and DV. The gain for alliance 2, in this situation, is very minor with BPS. DV, however, causes losses for alliance 2 when its competitor also uses DV.

Alliance1	Alliance2	Results of	RM only	ΔBPS	ΔDV
DAVN	DAVN	Alliance 1	2.23	0.04	0.12
		Alliance 2	1.54	0.05	-0.17
ProBP	ProBP	Alliance 1	1.94	0.06	0.08
		Alliance 2	1.63	0.02	-0.28

Figure 8. Both alliances use BPS or DV, high demand only.

D. Discussion of Differences between Alliances 1 and 2

We have seen that BPS with dual control is the only method that consistently helps both alliances. BPS works best for alliance 2 in all cases, while DV works best for alliance 1. This is because the network characteristics of the 2 alliances differ: code share average fares are lower for alliance 2 than for alliance 1, and average local fares are higher. Also, locals comprise the largest proportion of revenues for alliance 2, and code shares a smaller proportion, as compared with alliance 1. The lower code share fares combined with the higher local fares causes alliance 2 to experience spiral down (this occurs when the RM system undervalues seats and opens too many low fares for booking) when using DV. Alliance 1 does not experience a revenue decline because of different network characteristics: higher code share fares, lower local and connecting fares, and higher revenue share from code shares.

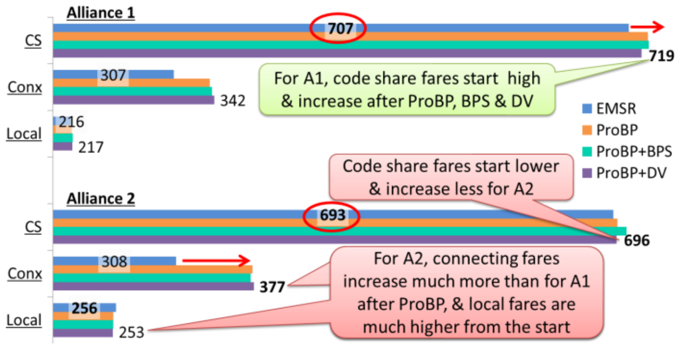

Figure 9. Average component fares for alliances 1 and 2.

Figure 9 illustrates the average component fares for alliances 1 and 2, and shows how the average fares change after the application of ProBP, and additionally BPS and DV. Code share fares increase more for alliance 1 than for alliance 2 when applying network RM and further alliance RM techniques. Connecting fares increase much more for alliance 2 than alliance 1, making that component more important for revenues. When DV is used by alliance 2, subtracting the already higher bid prices (resulting from the higher local fares of alliance 2) from the lower code share fares causes spiral down and harms revenues. Too many code share bookings are allowed at the expense of high-revenue local and connecting itineraries. This does not occur for alliance 1 because of its inherently higher code share fares and lower local fares.

V. DISCUSSION

A. Conclusions

In network E, which differs from US-based network A4 (used previously in this body of research in [15, 17, 18]) in terms of its physical structure, the properties of the types of flights (short-haul local and domestic connecting, along with long-haul international code share flights), and the benefits of BPS and DV, depend on the network characteristics to a larger degree. If code share flights are high-revenue relative to local itineraries, then BPS and DV improve revenues by raising bid prices (ProBP) or improving booking limits (DAVN). These benefits are more pronounced, and do not affect the revenues from the other traffic components, when code share itineraries comprise a larger proportion of revenues. If, however, the airlines in an alliance do not obtain the majority of revenues from code share traffic and the average code share fare is relatively low compared with the average local fare, then the potential danger of displacing a local passenger and causing revenue loss presents a problem for implementing BPS and DV.

The benefits from BPS and DV are present in most cases for alliances of differing structure, and they are larger if the competition is using less sophisticated RM methods. Additionally, if two competing alliances differ in their network characteristics such that one carries higher-revenue code share passengers and obtains a large amount of revenues from that component, then it could experience a larger benefit from BPS and DV to the detriment of the competitor. The results indicate that the frequent optimization of bid prices and booking limits may be particularly important in networks with a large difference in the lengths of their sets of flight legs, intense competition between two alliances, and fare structures that are prone to spiral down.

B. Future Research Directions

The two networks, A4 [15, 17, 18] and E, used in this body of research had specific network structures that resulted in somewhat different conclusions about the performance of BPS and DV. Expanding this research to networks of various other structures, and with other characteristics (e.g., the ratios of the average code share fare to local and connecting fares, and the proportions of revenue derived from the various traffic components) would help to compile a more holistic picture of the benefits of BPS and DV as a function of network structure.

Some literature on the alliance RM topic [23] has modeled the effects of partners choosing which bid prices to post for their own flight legs to maximize their own individual revenues, assuming that that the posted bid price is equal to the "transfer price", or the revenue that will receive from the operating partner if the itinerary is sold (revenues are not split according to a prorate agreement in this case). In addition to this "partner price" scenario, it is possible that some degree of bid price scaling (either up or down), in certain cases, would produce better availability decisions. A topic for further research is the effect on total alliance revenues of modifying bid prices for either mutual or individual benefit.

This paper was concerned with the combined alliance revenues and has assumed that the revenue resolution contracts are fixed and have been negotiated such that each partner views his share as a fair division. Another direction for future research is to examine the effects of different revenue resolution schemes or prorate agreements between alliance partners on the behavior and incentives of individual participants. Some airlines use fare adjustment (also called "marginal revenue optimization"), which decreases the fare inputs to the optimizer of lower class itineraries [11]. Some bid prices will be much lower if using fare adjustment. The performance of BPS may be significantly different if airlines are using this method. An idea suggested by [9] is that of bid price inference, where the lowest available fare on a flight leg is used as a pseudo-bid price. If actual bid prices are skewed downwards because of techniques like fare adjustment, then using lowest available fares may be a feasible alternative. Bid price inference is also applicable in the case when airlines do not have antitrust immunity to share bid prices, in which case communicating to each other the lowest available fares (which are public and would not require immunity) is still possible. Research on the success of BPS when using pseudo-bid prices, because it is unsuitable or infeasible to exchange actual bid prices, is another topic of interest.

We have seen that DAVN performs well as a network RM method because its composition of multiple heuristics makes it robust, and alone (without BPS) it produces large revenue gains even without frequent optimization. However, in the network E environment, applying BPS or DV to DAVN with time frame optimization resulted in large revenue losses. The work of [15] showed that a different, stricter implementation of BPS in DAVN resulted in better performance. Because DAVN comprises multiple heuristics, it may require airline-specific tweaking to obtain maximum performance. Further research into the adjustments needed to improve the performance of BPS in DAVN, and in what situations such adjustments are appropriate, would be of practical importance.

Also of practical importance to the industry is to test BPS and DV with different RM combinations that are not limited to the O-D control methods of DAVN and ProBP. We have tested various network and leg RM combinations in another network (A4), but it is important to continue this research for networks of varying structures as well, and confirm whether BPS with dual airline control remains the only method that consistently delivers revenue gains, though marginal in some cases. To the author's knowledge, most of the airlines in alliances are only just taking the first steps towards sharing bid price information and incorporating it into their code share availability decisions.

A major limitation of the models and methods proposed in the literature thus far is that they were tested on relatively small and simple hypothetical networks. Other than this research, the author is unaware of literature that also examined the performance of alliance RM in a competitive alliance environment. Once such alliance RM systems are in place as those proposed here and in the literature, data documenting the code share availability decisions and actual mechanisms used by airlines in real world competitive and network scenarios can be collected. The success of the various methods proposed here and elsewhere in the literature can then be validated according to real airline data.

C. Implications for Airlines

The results show that airlines in alliances who are considering sharing information about the value of seats on their flight legs should carefully consider their network structure and competitive environment before deciding to implement such methods into their availability control. If the network structure and competitive environment are conducive to obtaining large benefits from sharing bid price information, then large benefits will compound over time, outweighing the one time expense and technological complexity of implementation. Gains over $75 million per year can be obtained by a large airline that participates in an alliance such as American Airlines (assuming an approximate percent revenue gain of +0.30%) [14].

Although DV can produce the largest revenue gains for an alliance, particularly when code shares are a valuable and important part of the bookings and when the competitor uses less advanced RM, it is also more costly or difficult to implement because it requires modification of the RM system's valuation procedure. BPS, on the other hand, allows the RM system to function unmodified, and requires a real time test by both partners after own booking limits have been determined. Thus, BPS may be easier for airlines to implement. Although BPS does not always provide the highest revenue gains, the results in this paper showed that it is a more robust method for airlines whose network structures make them more susceptible to spiral down when they use DV (such as alliance 2). Whether

BPS or DV is a worthwhile investment is ultimately for airlines to decide.

In an industry with tiny profit margins and huge costs, small revenue benefits can make a big difference. This research has shown that alliance partners can cooperate and share their bid prices, whether they take the form of prorated expected values of empty seats on flight legs, or network displacement costs approximated as shadow prices from a deterministic LP, to attempt to improve the total alliance revenues. Through the use of bid price sharing for post-fact code share itinerary availability control, or by dynamically valuing the code share itinerary's revenue contribution in the network optimizer, airlines in alliances can improve their revenues and affirm the benefits of entering into code sharing agreements.

ACKNOWLEDGMENT

This research was supported by the PODS consortium's participant airlines. Thank you to Lufthansa, United, Swiss, SAS, Air Canada, Delta and KLM whose involvement in the research process and vocal participation in discussions about alliance revenue management during PODS conferences were particularly helpful. Without the PODS simulator research tool (developed at The Boeing Company) and improvements to its capability throughout this research by Craig Hopperstad, this research would not have been possible.

REFERENCES

[1] Belobaba, P. P. 1987. Air Travel Demand and Airline Seat Inventory Management.Ph.D. thesis, Massachusetts Institute of Technology, Flight Transportation.

[2] Belobaba, P. P. 1989. Application of probabilistic decision model to airline seat inventory control.Operations Research. 37(2) 182–197.

[3] Belobaba, P. P. 1992. The Revenue Enhancement Potential of Airline Yield Management Systems In: Proceedings of the 2nd International Conference on Advanced Software Technology for Air Transport, November 1992. 143-164.

[4] Belobaba, P. P. 2002. Airline Network Revenue Management: Recent Developments and State of the Practice. The Handbook of Airline Economics.The Aviation Weekly Group of the McGraw-Hill Companies, New York. Chapter 10, 141-156.

[5] Belobaba, P. P., A. R. Odoni, C. Barnhart. 2009. The Global Airline Industry. John Wiley & Sons, Ltd.

[6] Belobaba, P.P. 2010. Passenger Origin-Destination Simulator (PODS): Summary of Processes and Models. PODS Consortium. September, 2010.

[7] Boyd, A. E., 1998. Airline alliances. OR/MS Today 25(10), 28-31.

[8] Bratu, S. 1998. Network Value Concept in Airline Revenue Management.Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

[9] Darot, J. F. J. 2001. Revenue Management for Airline Alliances: Passenger Origin-Destination Simulation Analysis. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

[10] De la Torre, P. E. F. 1999. Airline Alliances: The Airline Perspective. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

[11] Fiig, T., Isler, K., Hopperstad, C., Belobaba, P., 2010. Optimization of mixed fare structures: Theory and applications. Journal of Revenue and Pricing Management 9, 152-170.

[12] Graf, M., Kimms, A., 2011. An option-based revenue management procedure for strategic airline alliances. European Journal of Operational Research 215(2), 459{469.

[13] Hopperstad, Craig. 2005. Passenger O/D Simulator PODS. Hopperstad Consulting, Seattle.v.10 edition, September, 2005.

[14] International Air Transport Association, 2014. Airport Intelligence Services Traffic Per Region Report for the 12 Months Ended December 2013. February, 2014. airport-is.com

[15] Jain, H. 2011. Alliance Revenue Management in Practice: Techniques and Simulation Analysis. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

[16] Lufthansa. 2011. The Lufthansa Group reports operating profit of EUR 578m after nine months. Available at http://presse.lufthansa.com/en/news-releases/singleview/archive/2011/october/27/article/2022.html

[17] Michel. A 2012. Airline alliances: Effects on total revenues of sharing estimated network value of unsold seats. Proceedings of International Conference on Reserch in Air Transportation. Berkeley, California. May 22-25, 2012.

[18] Michel, A. 2012. Airline Alliance Revenue Management:Improving Joint Revenues through Partner Sharing of Flight Leg Opportunity Costs Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

[19] Steer Davies Gleave, 2007. Competition impact of airline code-share agreements, p. 7, http://ec.europa.eu/competition/sectors/transport/reports/airlinecode share.pdf

[20] Topaloglu, H., 2012. A duality based approach for network revenue management in airline alliances. Journal of Revenue and Pricing Management (forthcoming with advance online publication).

[21] Vinod, B. 2005.Alliance Revenue Management.Journal of Revenue and Pricing Management. 4(1) 66-82.

[22] Williamson, E. L. 1992. Airline Network Seat Inventory Control: Methodologies and Revenue Impacts. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.

[23] Wright, C. P., H. Groenevelt, R.A. Shumsky. 2010. Dynamic Revenue Management in Airline Alliances. Transportation Science. 44(1) 15-3

[24] Wright, C. P., 2010. Modeling and analysis of revenue management in airline alliances. Ph.D. thesis, William E. Simon Graduate School of Business Administration at the University of Rochester.